arXiv:1704.08447v1 [stat.ME] 27 Apr 2017

Regression Type Models for Extremal Dependence

Linda Mhalla, Miguel de Carvalho, Valérie Chavez-Demoulin *

Abstract

We propose a vector generalized additive modeling framework for taking into account the effect of covariates on angular density functions in a multivariate extreme value context. The proposed methods are tailored for settings where the dependence between extreme values may change according to covariates. We devise a maximum penalized log-likelihood estimator, discuss details of the estimation procedure, and derive its consistency and asymptotic normality. The simulation study suggests that the proposed methods perform well in a wealth of simulation scenarios by accurately recovering the true covariate-adjusted angular density. Our empirical analysis reveals an interesting contrast between the dynamics governing the extremal dependence of stock market losses for major players in the banking and healthcare industries. Supplementary material for this article is available online.

Keywords: Angular density; Covariate-adjustment; Penalized log-likelihood; Regression splines; Statistics of multivariate extremes

*Linda Mhalla is PhD Candidate, Geneva School of Economics and Management (GSEM), Research Center for Statistics, Université de Genève, Switzerland (*linda.mhalla@unige.ch*). Miguel de Carvalho is Assistant Professor, School of Mathematics, University of Edinburgh, UK (*miguel.decarvalho@ed.ac.uk*). Valérie Chavez-Demoulin is Professor, Faculty of Business and Economics (HEC), Université de Lausanne, Switzerland (*valerie.chavez@unil.ch*).
The research was partially funded by FCT (Fundação para a Ciência e a Tecnologia, Portugal) through the project UID/MAT/00006/2013. The authors thank participants of Workshop 2017, EPFL, for discussions and comments. They also would like to thank Paul Embrechts for his constant encouragement.

1 INTRODUCTION

In this paper, we address an extension of the standard approach for modeling non-stationary univariate extremes to the multivariate setting. In the univariate context, the limiting distribution for the maximum of a sequence of independent and identically distributed random variables, derived by Fisher and Tippett (1928), is given by a generalized extreme value distribution characterized by three parameters: μ (location), σ (scale), and ξ (shape). To take into account the effect of a vector of covariates $\mathbf{x}$, one can let these parameters depend on $\mathbf{x}$ and the resulting generalized extreme value distribution takes the form

$$G_{(\mu_{\mathbf{x}},\sigma_{\mathbf{x}},\xi_{\mathbf{x}})}(y) = \exp\left[-\left\{1+\xi_{\mathbf{x}}\left(\frac{y-\mu_{\mathbf{x}}}{\sigma_{\mathbf{x}}}\right)\right\}_{+}^{-1/\xi_{\mathbf{x}}}\right], \tag{1}$$

where $(a)_+ = \max\{0, a\}$, see Coles (2001, ch. 6), Pauli and Coles (2001), Chavez-Demoulin and Davison (2005), Yee and Stephenson (2007), Wang and Tsai (2009), Eastoe and Tawn (2009), and Chavez-Demoulin et al. (2015), for related approaches.

In the multivariate context, consider $\mathbf{Y}_i = (Y_{i,1}, \ldots, Y_{i,d})^{\mathrm{T}}$ independent and identically distributed random vectors with joint distribution F, and unit Fréchet marginal distribution functions $F_j(y) = \exp(-1/y)$, for $y > 0$. Pickands' representation theorem (Coles, 2001, Theorem 8.1) states that the law of the standardized componentwise maxima, $\mathbf{M}_n = n^{-1}\max\{\mathbf{Y}_1, \ldots, \mathbf{Y}_n\}$, converges in distribution to a multivariate extreme value distribution, $G_H(\mathbf{y}) = \exp\{-V_H(\mathbf{y})\}$, with

$$V_H(\mathbf{y}) = \int_{S_d} \max\left(\frac{w_1}{y_1}, \ldots, \frac{w_d}{y_d}\right) \mathrm{d}H(\mathbf{w}). \tag{2}$$

Here H is the so-called angular measure, that is, a positive finite measure on the unit simplex $S_d = \left\{(w_1, \ldots, w_d) \in [0,\infty)^d : w_1 + \cdots + w_d = 1\right\}$ that needs to obey

$$\int_{S_d} w_j \, \mathrm{d}H(\mathbf{w}) = 1, \quad j = 1, \ldots, d. \tag{3}$$

The function $V(\mathbf{y}) \equiv V_H(\mathbf{y})$, is the so-called exponent measure, and is continuous, convex, and homogeneous of order -1, i.e., $V(t\mathbf{y}) = t^{-1}V(\mathbf{y})$, for all $t > 0$.

The class of limiting distributions of multivariate extreme values yields an infinite number of possible parametric representations (Coles, 2001, ch. 8), as the validity of a multivariate extreme value distribution is conditional on its angular measure H verifying the moment constraint (3). Therefore, most literature has focused on the estimation of the extremal dependence structures described by spectral measures or equivalently angular densities (Boldi and Davison, 2007; Einmahl et al., 2009; de Carvalho et al., 2013; Sabourin and Naveau, 2014; Hanson et al., 2017). Related

quantities such as the Pickands dependence function (Pickands, 1981) and the stable tail dependence function (Huang, 1992; Drees and Kaufmann, 1998) were investigated by many authors (Einmahl et al., 2006; Gudendorf and Segers, 2012; Wadsworth and Tawn, 2013; Marcon et al., 2016). A wide variety of parametric models for the spectral density that allow flexible dependence structures were proposed (Kotz and Nadarajah, 2000, sec. 3.4).

However, few papers were able to satisfactorily address the challenging but incredibly relevant setting of modeling nonstationarity at joint extreme levels. Some exceptions include de Carvalho and Davison (2014), who proposed a nonparametric approach where a family of spectral densities is constructed using exponential tilting. Castro and de Carvalho (2017) developed an extension of this approach based on covariate-varying spectral densities. These approaches are however limited to replicated one-way ANOVA types of settings. de Carvalho (2016) advocated the use of covariate-adjusted angular densities, and Escobar-Bach et al. (2016) discussed estimation—in the bivariate and covariate-dependent framework—of the Pickands dependence function based on local estimation with minimum density power divergence criterion. Finally, Mhalla et al. (2017) constructed, in a nonparametric framework, smooth models for predictor-dependent Pickands dependence functions based on generalized additive models.

Our approach is based on a non-linear model for covariate-varying extremal dependences. Specifically, we develop a vector generalized additive model which flexibly allows the extremal dependence to change with a set of covariates, but—keeping in mind that extreme values are scarce—it borrows strength from a parametric assumption. In other words, the goal is to develop a regression model for the extremal dependence through the parametric specification of an extremal dependence structure and then to model the parameters of such a structure through a vector generalized additive model (VGAM) (Yee and Wild, 1996; Yee, 2015). One major advantage over existing methods is that our model may be used for handling an arbitrary number of dimensions and covariates as it takes advantage of the resilience of the GAM paradigm, and it is straightforward to implement as illustrated in the R code (R Development Core Team, 2016) in the Supplementary Material.

The remainder of this paper is organized as follows. In Section 2 we introduce the proposed model for covariate-adjusted extremal dependences. In Section 3 we develop our penalized likelihood approach and give details on the asymptotic properties of our estimator. In Section 4 we assess the performance of the proposed methods. An application to stock market data is given in Section 5. We close the paper in Section 6 with a discussion. The Supplementary Material contains a `knitR`-produced file (Xie, 2015) with R code to implement the proposed methodologies.

2 FLEXIBLE COVARIATE-ADJUSTED ANGULAR DENSITIES

2.1 Statistics of Multivariate Extremes: Preparations and Background

The functions H and V in (2) can be used to describe the structure of dependence between the extremes, as in the case of independence between the extremes where $V(\mathbf{y}) = \sum_{j=1}^{d} 1/y_d$ and in the case of perfect extremal dependence where $V(\mathbf{y}) = \max\{1/y_1, \ldots, 1/y_d\}$. As a consequence, if H is differentiable with angular density denoted h, the more mass around the barycenter of S_d, $(d^{-1}, \ldots, d^{-1})$, the higher the level of extremal dependence. Further insight into these measures may be obtained by considering the point process $P_n = \{n^{-1}\mathbf{Y}_i : i = 1, \ldots, n\}$. Following de Haan and Resnick (1977) and Resnick (1987, sec. 5.3), as $n \to \infty$, P_n converges to a non-homogeneous Poisson point process P defined on $[\mathbf{0}, \boldsymbol{\infty}) \setminus \{\mathbf{0}\}$ with a mean measure μ that verifies

$$\mu(A_{\mathbf{y}}) = V(\mathbf{y}),$$

where $A_{\mathbf{y}} = \mathbb{R}^d \setminus ([-\infty, y_1] \times \cdots \times [-\infty, y_d])$.

There are two representations of the intensity measure of the limiting Poisson point process P which will be handy for our purposes. First, it holds that

$$\mu(\mathrm{d}\mathbf{y}) = -V_{1:d}(\mathbf{y})\ \mathrm{d}\mathbf{y}, \tag{4}$$

with $V_{1:d}$ being the derivative of V with respect to all its arguments (Resnick, 1987, sec. 5.4). Second, another useful factorization of the intensity measure $\mu(\mathrm{d}\mathbf{y})$, called the spectral decomposition, can be obtained using the following decomposition of the random variable $\mathbf{Y} = (Y_1, \ldots, Y_d)^{\mathrm{T}}$ into radial and polar coordinates,

$$(R, \mathbf{W}) = \left(\|\mathbf{Y}\|, \frac{\mathbf{Y}}{\|\mathbf{Y}\|} \right), \tag{5}$$

where $\|\cdot\|$ denotes the L_1-norm. Indeed, it can be shown that (Beirlant et al., 2004, sec. 8.2.3) the limiting intensity measure factorizes across radial and angular components as follows

$$\mu(\mathrm{d}\mathbf{y}) = \mu(\mathrm{d}r \times \mathrm{d}\mathbf{w}) = \frac{\mathrm{d}r}{r^2}\, \mathrm{d}H(\mathbf{w}).$$

The spectral decomposition (5) allows the separation of the marginal and the dependence parts in the multivariate extreme value distribution G_H with the margins being unit Fréchet and the dependence structure being described by the angular measure H.

The inference approach which we build on in this paper is that developed by Coles and Tawn (1991) and based on threshold excesses; see Huser et al. (2016) for a detailed review of likelihood estimators for multivariate extremes. The set of extreme events is defined as the set of observations

with radial components exceeding a high fixed threshold, i.e., the observations belonging to the extreme set,

$$E_{\mathbf{r}} = \left\{ (y_1, \ldots, y_d) \in (0, \infty)^d : \sum_{j=1}^{d} \frac{y_j}{r_j} > 1 \right\},$$

with $\mathbf{r} = (r_1, \ldots, r_d)$ a large threshold vector. Since the points $n^{-1}\mathbf{Y}_i$ are mapped to the origin for non-extreme observations, the threshold $\mathbf{r}$ needs to be sufficiently large for the Poisson approximation to hold. Note that, $\mathbf{Y}_i \in E_{\mathbf{r}}$, if and only if,

$$R_i = \|\mathbf{Y}_i\| > \left(\sum_{j=1}^{d} \frac{W_{i,j}}{r_j} \right)^{-1}, \quad \text{where } W_{i,j} = \frac{Y_{i,j}}{R_i}.$$

Hence, the expected number of points of the Poisson process P located in the extreme region $E_{\mathbf{r}}$ is

$$\begin{aligned} \mu(E_{\mathbf{r}}) &= \int_{S_d} \int_{\left(\sum_{j=1}^{d} \frac{w_j}{r_j}\right)^{-1}}^{\infty} \frac{\mathrm{d}r}{r^2} \, \mathrm{d}H(\mathbf{w}) \\ &= \int_{S_d} \left(\sum_{j=1}^{d} \frac{w_j}{r_j} \right) \mathrm{d}H(\mathbf{w}) \\ &= \sum_{j=1}^{d} \frac{1}{r_j} \int_{S_d} w_j \, \mathrm{d}H(\mathbf{w}) = \sum_{j=1}^{d} \frac{1}{r_j}. \end{aligned} \tag{6}$$

Now, we can explicitly formulate the Poisson log-likelihood over the set E_r,

$$\ell_{E_{\mathbf{r}}}(\boldsymbol{\theta}) = -\mu(E_{\mathbf{r}}) + \sum_{i=1}^{n_{\mathbf{r}}} \log \left\{ \mu(\mathrm{d}R_i \times \mathrm{d}\mathbf{W}_i) \right\}, \tag{7}$$

where $\boldsymbol{\theta}$ represents the vector of parameters of the measure μ and $n_{\mathbf{r}}$ the number of reindexed observations in the extreme set $E_{\mathbf{r}}$. Using (6), the first term in (7) can be omitted when maximizing the Poisson log-likelihood which, using (4), boils down to

$$\ell_{E_{\mathbf{r}}}(\boldsymbol{\theta}) \equiv \sum_{i=1}^{n_{\mathbf{r}}} \log \left\{ -V_{1:d}(\mathbf{Y}_i; \boldsymbol{\theta}) \right\}. \tag{8}$$

Thanks to the differentiability of the exponent measure V, and the support of the angular measure H in the unit simplex S_d, we can use a result by Coles and Tawn (1991, Theorem 1) that relates the angular density to the exponent measure via

$$V_{1:d}(\mathbf{y}) = -\|\mathbf{y}\|^{-(d+1)} h\left(\frac{y_1}{\|\mathbf{y}\|}, \ldots, \frac{y_d}{\|\mathbf{y}\|}; \boldsymbol{\theta} \right),$$

and reformulate the log-likelihood (8) as follows

$$\begin{aligned} \ell_{E_{\mathbf{r}}}(\boldsymbol{\theta}) &\equiv -(d+1) \sum_{i=1}^{n_{\mathbf{r}}} \log \|\mathbf{Y}_i\| + \sum_{i=1}^{n_{\mathbf{r}}} \log \left\{ h\left(\frac{Y_{i,1}}{\|\mathbf{Y}_i\|}, \ldots, \frac{Y_{i,d}}{\|\mathbf{Y}_i\|}; \boldsymbol{\theta} \right) \right\} \\ &= \sum_{i=1}^{n_{\mathbf{r}}} \ell_{E_{\mathbf{r}}}(\mathbf{Y}_i, \boldsymbol{\theta}). \end{aligned} \tag{9}$$

2.2 Vector Generalized Additive Models for Covariate-Adjusted Angular Densities

Our starting point for modeling is an extension of (1) to the multivariate setting. Whereas the model in (1) is based on indexing the parameters of the univariate extreme value distribution with a regressor, here we index the parameter (H) of a multivariate extreme value distribution (G_H), with a regressor $\mathbf{x} = (x_1, \ldots, x_q)^{\mathrm{T}} \in \mathcal{X} \subset \mathbb{R}^q$. Our target object of interest is thus given by a family of covariate-adjusted angular measures $H_{\mathbf{x}}$ obeying

$$\int_{S_d} w_j \mathrm{d}H_{\mathbf{x}}(\mathbf{w}) = 1, \quad j = 1, \ldots, d.$$

Of particular interest is the setting where $H_{\mathbf{x}}$ is differentiable, in which case the covariate-adjusted angular density can be defined as $h_{\mathbf{x}}(\mathbf{w}) = \mathrm{d}H_{\mathbf{x}}/\mathrm{d}\mathbf{w}$. This yields a corresponding family of covariate-indexed multivariate extreme value distributions

$$G_{\mathbf{x}}(\mathbf{y}) = \exp\left\{ -\int_{S_d} \max\left(\frac{w_1}{y_1}, \ldots, \frac{w_d}{y_d}\right) \mathrm{d}H_{\mathbf{x}}(\mathbf{w}) \right\}.$$

Some parametric models are used below to illustrate the concept of covariate-adjusted angular densities, and we focus on the bivariate setting for the sake of illustrating ideas. To develop insight and intuition on these models, see Figure 1.

Example 1 (Logistic angular surface)**.** Let

$$h_{\mathbf{x}}(w) = (1/\alpha_{\mathbf{x}} - 1)\left\{w(1-w)\right\}^{-1-1/\alpha_{\mathbf{x}}}\left\{w^{-1/\alpha_{\mathbf{x}}} + (1-w)^{-1/\alpha_{\mathbf{x}}}\right\}^{\alpha_{\mathbf{x}}-2}, \quad w \in (0,1), \tag{10}$$

with $\alpha : \mathcal{X} \subset \mathbb{R}^q \to (0,1]$. In Figure 1 a) we represent the case $\alpha_x = x$, with $x \in \mathcal{X} = [0.2, 0.9]$, which corresponds to be transitioning between a case of relatively high extremal dependence (lower values of x) to one of asymptotic independence (higher values of x).

Example 2 (Dirichlet angular surface)**.** Let

$$h_{\mathbf{x}}(w) = \frac{\alpha_{\mathbf{x}}\beta_{\mathbf{x}}\Gamma(\alpha_{\mathbf{x}} + \beta_{\mathbf{x}} + 1)(\alpha_{\mathbf{x}} w)^{\alpha_{\mathbf{x}}-1}\{\beta_{\mathbf{x}}(1-w)\}^{\beta_{\mathbf{x}}-1}}{\Gamma(\alpha_{\mathbf{x}})\Gamma(\beta_{\mathbf{x}})\{\alpha_{\mathbf{x}} w + \beta_{\mathbf{x}}(1-w)\}^{\alpha_{\mathbf{x}}+\beta_{\mathbf{x}}+1}}, \quad w \in (0,1),$$

with $\alpha : \mathcal{X} \subset \mathbb{R}^q \to (0,\infty)$ and $\beta : \mathcal{X} \subset \mathbb{R}^q \to (0,\infty)$. In Figure 1 b) we consider the case $\alpha_x = \exp(x)$ and $\beta_x = x^2$, with $x \in [0.9, 3]$. Note the different schemes of extremal dependence induced by the different values of the covariate x as well as the asymmetry of the spectral surface underlying this model.

Example 3 (Hüsler–Reiss angular surface)**.** Let

$$h_{\mathbf{x}}(w) = \frac{1}{\lambda_{\mathbf{x}} w^2 (1-w)(2\pi)^{1/2}} \exp\left\{ - \frac{\left[\log\{(1-w)/w\} + 2\lambda_{\mathbf{x}}^2\right]^2}{8\lambda_{\mathbf{x}}^2} \right\}, \quad w \in (0,1),$$

where $\lambda : \mathcal{X} \subset \mathbb{R}^q \to (0, \infty)$. In Figure 1 c) we consider the case $\lambda_x = \exp(x)$, with $x \in [0.1, 2]$. Under this specification, lower values of x correspond to lower levels of extremal dependence, whereas higher values of x correspond to higher levels of extremal dependence.

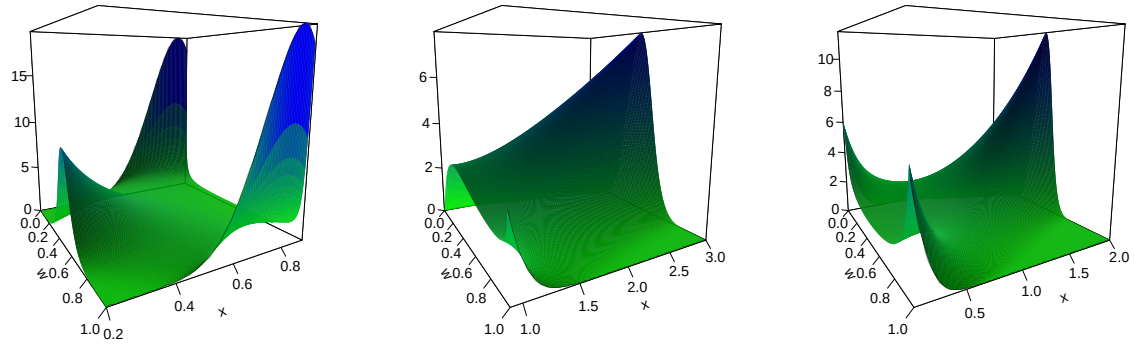

(a) (b) (c)

Figure 1: Covariate-adjusted angular densities of logistic (a), Dirichlet (b), and Hüsler–Reiss (c) models, corresponding respectively to the specifications in Examples 1, 2, and 3.

The previous parametric models provide some examples of covariate-adjusted angular surfaces $h_{\mathbf{x}}$. But how can we learn about $h_{\mathbf{x}}$ from the data? Suppose we observe the regression data $\{(\mathbf{x}_i, \mathbf{Y}_i)\}_{i=1}^n$, with $(\mathbf{x}_i, \mathbf{Y}_i) \in \mathbb{R}^q \times \mathbb{R}^d$, and where we assume that $\mathbf{Y}_i = (Y_{i,1}, \ldots, Y_{i,d})^{\mathrm{T}}$ are independent random vectors with unit Fréchet marginal distributions. Using a similar approach as in Section 2.1, we convert the raw sample into a pseudo-sample of cardinality $n_{\mathbf{r}}$,

$$\{(\mathbf{x}_i, \mathbf{Y}_i) : \mathbf{Y}_i \in E_{\mathbf{r}}\},$$

and use the latter reindexed data to learn about $h_{\mathbf{x}}$. We model $h_{\mathbf{x}}(\mathbf{w})$ using $h(\mathbf{w}; \boldsymbol{\theta}_{\mathbf{x}})$, where the parameter underlying the dependence structure

$$\boldsymbol{\theta}_{\mathbf{x}} = (\theta_{1\mathbf{x}}, \ldots, \theta_{p\mathbf{x}})^{\mathrm{T}}, \quad \mathbf{x} = (x_1, \ldots, x_q)^{\mathrm{T}} \in \mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_q \subseteq \mathbb{R}^q,$$

is specified through a vector generalized additive model (VGAM) (Yee and Wild, 1996). Specifically, we model $h_{\mathbf{x}}(\mathbf{w})$ using a fixed parametric extremal dependence structure $h(\mathbf{w}; \boldsymbol{\theta}_{\mathbf{x}})$ (say, the Hüsler–Reiss model in Example 3, in which case $\theta_x \equiv \lambda_x = \exp(x)$, $p = 1$, $q = 1$, and $\mathcal{X} = \mathbb{R}$). Then, to

learn about $\boldsymbol{\theta}_{\mathbf{x}}$ from the pseudo-sample, we use a vector generalized additive model which takes the form

$$\mathbf{g}(\boldsymbol{\theta}_{\mathbf{x}}) = \sum_{k=1}^{q} \mathbf{f}_k(x_k). \tag{11}$$

Here,

- $\mathbf{g} = (g_1, \ldots, g_p)^{\mathrm{T}}$ is a p-vector of link functions, and
- $\mathbf{f}_k = (\mathbf{f}_{1,k}, \ldots, \mathbf{f}_{p,k})^{\mathrm{T}}$, where $\mathbf{f}_{j,k} = (f_{j,k}(x_{1k}), \ldots, f_{j,k}(x_{n_{\mathbf{r}}k}))$, and $f_{j,k} : \mathcal{X}_k \to \mathbb{R}$ are smooth functions supported on $\mathcal{X}_k$, for $k = 1, \ldots, q$ and $j = 1, \ldots, p$. The vectors of smooth functions $\mathbf{f}_k$ admit a quadratic penalty representation (see Section 3.1 for details).

The vector of parameters to be estimated in the VGAM (11) is $\boldsymbol{\psi} = (\mathbf{f}_1, \ldots, \mathbf{f}_q) \in \Psi \subseteq \mathbb{R}^{p \times (qn_{\mathbf{r}})}$.

The specification in (11) allows to fit simultaneously ordinary GAMs in each component of the vector $\boldsymbol{\theta}_{\mathbf{x}}$, avoiding hence any non-orthogonality related issues that could arise if the p components were to be treated separately (Chavez-Demoulin and Davison, 2005). It should be noted that a slighter complex version of the VGAM (11) can be implemented to tune the effects of the covariates on each component, for example impose the same smooth effect of a covariate on different components, restrict the effect of a covariate on some components, etc.

3 INFERENCE AND ASYMPTOTIC PROPERTIES

The log-likelihood (9) is now written as

$$\ell(\boldsymbol{\psi}) := \sum_{i=1}^{n_{\mathbf{r}}} \ell(\mathbf{Y}_i; \boldsymbol{\psi}) = \sum_{i=1}^{n_{\mathbf{r}}} \ell_{E_{\mathbf{r}}} \left\{ \mathbf{Y}_i, \mathbf{g}(\boldsymbol{\theta}_{\mathbf{x}_i}) \right\},$$

where $\boldsymbol{\psi} = (\mathbf{f}_1, \ldots, \mathbf{f}_q) \in \Psi \subseteq \mathbb{R}^{p \times (qn_{\mathbf{r}})}$ and $\mathbf{g}(\boldsymbol{\theta}_{\mathbf{x}_i})$ is defined as in (11). Incorporating a covariate-dependence in the extremal dependence model through a non-linear smooth model adds a huge flexibility in the modeling of the dependence parameter $\boldsymbol{\theta}_{\mathbf{x}}$. The price to pay for this flexibility is reflected in the estimation procedure. As a matter of fact, the estimation of $\boldsymbol{\theta}_{\mathbf{x}}$ is performed by maximizing the penalized log-likelihood

$$\ell(\boldsymbol{\psi}, \boldsymbol{\gamma}) = \ell(\boldsymbol{\psi}) - \sum_{j=1}^{p} \sum_{k=1}^{q} \gamma_{j,k} \int_{\mathcal{X}_k} \{f''_{j,k}(x)\}^2 \, \mathrm{d}x. \tag{12}$$

The penalty term in (12) controls the wiggliness of the smooth functions and their fidelity to the data through the smoothing parameters $(\boldsymbol{\gamma})_{j,k} = \gamma_{j,k}$, with larger values leading to smoother fitted curves $\widehat{f}_{j,k}$.

The maximization of the penalized log-likelihood (12) is performed using a Newton–Raphson algorithm based on a vector backfitting algorithm. In the same train of thought as the usual backfitting algorithm (Hastie and Tibshirani, 1990), the vector backfitting algorithm fits at each step of the Newton–Raphson algorithm, a vector additive model to a set of pseudo-responses. We detail the fitting procedure in the following section.

3.1 Fitting Algorithm

We suppose that log-likelihood contributions from independent observations $\{(\mathbf{x}_i, \mathbf{Y}_i)\}_{i=1}^{n_\mathbf{r}}$ depend on the vector of parameters $\big(\mathbf{g}(\boldsymbol{\theta}_{\mathbf{x}_1}), \ldots, \mathbf{g}(\boldsymbol{\theta}_{\mathbf{x}_{n_\mathbf{r}}})\big)$ which is a function of q matrices of smooth curves $(\mathbf{f}_k)_{j,i} = f_{j,k}(x_{ik})$, $i = 1, \ldots, n_\mathbf{r}$, for $j = 1, \ldots, p$ and $k = 1, \ldots, q$. These smooth curves are encompassed in the parameter of interest $\boldsymbol{\psi} \in \Psi \subseteq \mathbb{R}^{p \times (qn_\mathbf{r})}$. We give the details of a Newton–Raphson iteration which updates current values $\boldsymbol{\psi}^{(a)} = (\mathbf{f}_1^{(a)}, \ldots, \mathbf{f}_q^{(a)})$ to $\boldsymbol{\psi}^{(a+1)} = (\mathbf{f}_1^{(a+1)}, \ldots, \mathbf{f}_q^{(a+1)})$. We set $\mathbf{x}_i = (x_{i1}, \ldots, x_{iq})^\mathrm{T}$ and define for $\boldsymbol{\psi} = (\mathbf{f}_1, \ldots, \mathbf{f}_q)$, the vector $(\boldsymbol{\psi}_k)_i = (\mathbf{f}_k)_i = ((\mathbf{f}_k)_{1,i}, \ldots, (\mathbf{f}_k)_{p,i})^\mathrm{T}$. The dependence structure parameter is now denoted by

$$\boldsymbol{\theta}_{\mathbf{x}_i}^{(a)} = \sum_{k=1}^{q} (\boldsymbol{\psi}_k^{(a)})_i, \quad i = 1, \ldots, n_\mathbf{r}.$$

By differentiating the penalized log-likelihood (12) with respect to $\boldsymbol{\psi}$ and considering the ath iteration of the Newton–Raphson algorithm, we find that

$$\boldsymbol{\psi}^{(a+1)} = \underset{\boldsymbol{\psi}=(\mathbf{f}_1,\ldots,\mathbf{f}_q)}{\arg\min} \sum_{i=1}^{n_\mathbf{r}} \left\{ \mathbf{z}_i^{(a)} - \sum_{k=1}^{q} (\boldsymbol{\psi}_k)_i \right\}^\mathrm{T} \boldsymbol{\Sigma}_i^{(a)^{-1}} \left\{ \mathbf{z}_i^{(a)} - \sum_{k=1}^{q} (\boldsymbol{\psi}_k)_i \right\} + \sum_{j=1}^{p} \sum_{k=1}^{q} \gamma_{j,k} \int_{\mathfrak{X}_k} \{f_{j,k}''(x)\}^2 \, \mathrm{d}x, \tag{13}$$

where

$$\mathbf{z}_i^{(a)} = \boldsymbol{\theta}_{\mathbf{x}_i}^{(a)} + \{\boldsymbol{\Sigma}_i^{(a)}\}^{-1} \mathbf{u}_i^{(a)}, \quad \mathbf{u}_i^{(a)} = \frac{\partial \ell_{E_\mathbf{r}} \{\mathbf{Y}_i, \mathbf{g}(\boldsymbol{\theta})\}}{\partial \boldsymbol{\theta}} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{\mathbf{x}_i}^{(a)}}, \quad \boldsymbol{\Sigma}_i^{(a)} = \frac{-\partial^2 \ell_{E_\mathbf{r}} \{\mathbf{Y}_i, \mathbf{g}(\boldsymbol{\theta})\}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^\mathrm{T}} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{\mathbf{x}_i}^{(a)}}.$$

In (13), $\boldsymbol{\psi}^{(a+1)}$ is the solution to the following vector additive model

$$\mathbf{z}_i^{(a)} = \sum_{k=1}^{q} (\boldsymbol{\psi}_k)_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim (\mathbf{0}, \boldsymbol{\Sigma}_i^{(a)}), \quad i = 1, \ldots, n_\mathbf{r}, \tag{14}$$

which is fitted using a vector backfitting algorithm (Yee, 2015) resulting in the following estimates

$$\boldsymbol{\psi}_k^{(a+1)} = \mathbf{f}_k^{(a+1)} = \mathbf{A}^{(a)}(\boldsymbol{\gamma}_k) \left\{ \mathbf{z}^{(a)} - \sum_{s \neq k} \mathbf{f}_s^{(a+1)} \right\}, \quad k = 1, \ldots, q,$$

where

$$\mathbf{z}^{(a)} = (\mathbf{z}_1^{(a)^\mathrm{T}}, \ldots, \mathbf{z}_{n_\mathbf{r}}^{(a)^\mathrm{T}})^\mathrm{T}, \quad \mathbf{A}^{(a)}(\boldsymbol{\gamma}_k) = (\mathbf{I}_{n_\mathbf{r} p} + \boldsymbol{\Sigma}^{(a)} \mathbf{K}_k)^{-1}.$$

The $n_{\mathbf{r}}p \times n_{\mathbf{r}}p$ matrices $\boldsymbol{\Sigma}^{(a)}$ and $\mathbf{K}_k$ are defined as follows

$$\boldsymbol{\Sigma}^{(a)} = \begin{pmatrix} \boldsymbol{\Sigma}_1^{(a)} & & \\ & \ddots & \\ & & \boldsymbol{\Sigma}_{n_{\mathbf{r}}}^{(a)} \end{pmatrix}, \quad \mathbf{K}_k = \mathbf{Q}_k \mathbf{T}_k^{-1} \mathbf{Q}_k \otimes \mathrm{diag}(\gamma_{1,k}, \ldots, \gamma_{p,k}),$$

where $\mathbf{Q}_k$ is a banded $n_{\mathbf{r}} \times (n_{\mathbf{r}} - 2)$ matrix, $\mathbf{T}_k$ is a symmetric tridiagonal matrix of order $n_{\mathbf{r}} - 2$, and $\otimes$ denotes the Kronecker product. The matrices $\mathbf{Q}_k$ and $\mathbf{T}_k$ depend only on the observed values of the covariates x_{ik}, for $i = 1, \ldots, n_{\mathbf{r}}$. When the information matrix $\boldsymbol{\Sigma}_i^{(a)}$ fails to be positive-definite, a Fisher scoring algorithm is implemented instead of the Newton–Raphson algorithm and consists of replacing $\boldsymbol{\Sigma}_i^{(a)}$ by its expectation $\mathrm{E}[-\partial^2 \ell_{E_{\mathbf{r}}}\{\mathbf{Y}, \mathbf{g}(\boldsymbol{\theta})\}/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}^{\mathrm{T}}]$ evaluated at $\boldsymbol{\theta} = \boldsymbol{\theta}_{\mathbf{x}_i}^{(a)}$.
The quadratic penalty representation assumption on $\mathbf{f}_k$, for $k = 1, \ldots, q$, can now be defined as

$$\sum_{j=1}^{p} \gamma_{j,k} \int_{\mathcal{X}_k} \{f_{j,k}''(x)\}^2 \, \mathrm{d}x = ((\mathbf{f}_k)_1^{\mathrm{T}}, \ldots, (\mathbf{f}_k)_{n_{\mathbf{r}}}^{\mathrm{T}}) \mathbf{K}_k ((\mathbf{f}_k)_1^{\mathrm{T}}, \ldots, (\mathbf{f}_k)_{n_{\mathbf{r}}}^{\mathrm{T}})^{\mathrm{T}}.$$

Finally, the vector of parameters $\boldsymbol{\psi}^{(a)}$ is updated by

$$\boldsymbol{\psi}^{(a+1)} = (\mathbf{f}_1^{(a+1)}, \ldots, \mathbf{f}_q^{(a+1)}).$$

The algorithm stops when the change in the resulting dependence structure parameters between two successive iterations is sufficiently small. The penalized maximum log-likelihood estimator (PMLE) is then defined as

$$\widehat{\boldsymbol{\psi}} = \arg \max_{\boldsymbol{\psi} \in \Psi} \ell(\boldsymbol{\psi}, \boldsymbol{\gamma}). \tag{15}$$

The corresponding plug-in PMLE of the conditional angular density is defined as

$$\widehat{h}_{\mathbf{x}}(\mathbf{w}) \equiv h(\mathbf{w}; \widehat{\boldsymbol{\theta}}_{\mathbf{x}}), \tag{16}$$

where $\widehat{\boldsymbol{\theta}}_{\mathbf{x}} = \mathbf{g}^{-1}\{\sum_{k=1}^{q} \widehat{\mathbf{f}}_k(x_k)\}$.

3.2 Selection of the Smoothing Parameters

At the ath iteration of the Newton–Raphson algorithm, the smoothing parameters $\boldsymbol{\gamma}_k = (\gamma_{1,k}, \ldots, \gamma_{p,k})$ related to the vector of smooth curves $\mathbf{f}_k^{(a+1)}$, for $k = 1, \ldots, q$, are chosen by minimizing the generalized cross validation (GCV) criterion

$$\mathrm{GCV}^{(a)} = \frac{1}{n_{\mathbf{r}}} \sum_{i=1}^{n_{\mathbf{r}}} \frac{\{\mathbf{z}_i^{(a)} - \sum_{s \neq k} (\mathbf{f}_s^{(a+1)})_i\}^{\mathrm{T}} {\boldsymbol{\Sigma}_i^{(a)}}^{-1} \{\mathbf{z}_i^{(a)} - \sum_{s \neq k} (\mathbf{f}_s^{(a+1)})_i\}}{n_{\mathbf{r}}^{-1} \mathrm{trace}\{\mathbf{I}_{n_{\mathbf{r}}p} - \mathbf{A}^{(a)}(\boldsymbol{\gamma}_k)\}}.$$

The degrees of freedom of the smooth curve $f_{k,j}^{(a+1)}$ for $k = 1, \ldots, q$ and $j = 1, \ldots, p$, are defined as

$$\mathrm{df}_{(k,j)}^{(a)} = \mathrm{trace}\left\{\mathrm{diag}(\mathbf{1}_{n_\mathbf{r}} \otimes \mathbf{e}_j)\mathbf{A}^{(a)}(\boldsymbol{\gamma}_k)\mathrm{diag}(\mathbf{1}_{n_\mathbf{r}} \otimes \mathbf{e}_j)\right\},$$

where $\mathbf{e}_j = (0, \ldots, 1, 0, \ldots, 0)^\mathrm{T} \in \mathbb{R}^{n_\mathbf{r}}$ and $\mathbf{1}_{n_\mathbf{r}} = (1, \ldots, 1)$. The degrees of freedom of the model (11) are defined as the final $\mathrm{df}_{(k,j)}^{(a)}$ obtained at convergence.

The vector backfitting algorithm performed to fit the vector additive model (14), uses vector smoothing splines with components that can be shown to be natural cubic splines and can hence be written as linear combinations of B-splines. The smooth curves $f_{k,j}$ for $k = 1, \ldots, q$ and $j = 1, \ldots, p$, can then be written as

$$f_{k,j}(x) = \sum_{i=1}^{n_{k,j}^*+2} b_{k,j_i} B_i(x),$$

where $n_{k,j}^*$ is the number of fixed knots associated to $f_{k,j}$, b_{k,j_i}'s are coefficients, and B_i are B-spline basis functions. Therefore, the estimation of the vector of parameters $\boldsymbol{\psi} \in \mathbb{R}^{p\times(qn_\mathbf{r})}$ is reduced to the estimation of the $\sum_{j=1}^{p}\sum_{k=1}^{q} n_{k,j}^*$ B-splines coefficients. In practice, we use the B-spline estimation as it is expected to be more numerically stable (Yee, 2015). Further details on the estimation of VGAMs as well as on the consistency and convergence of the vector backfitting algorithm can be found in Yee (2015, sec. 4.3.2).

3.3 Large Sample Properties

In this section we derive the consistency and asymptotic normality of the PMLE $\widehat{\boldsymbol{\psi}}$ defined in (15). For $j = 1, \ldots, p$, we define

$$\boldsymbol{\gamma}^j = (\gamma_{j,1}, \ldots, \gamma_{j,q}), \quad \boldsymbol{\psi}^j = ((\mathbf{f}_1)_{j,1}, \ldots, (\mathbf{f}_1)_{j,n_\mathbf{r}}, \ldots, (\mathbf{f}_q)_{j,1}, \ldots, (\mathbf{f}_q)_{j,n_\mathbf{r}})^\mathrm{T}.$$

Hence, $\boldsymbol{\psi} = ({\boldsymbol{\psi}^1}^\mathrm{T}, \ldots, {\boldsymbol{\psi}^p}^\mathrm{T})^\mathrm{T} \in \boldsymbol{\Psi} \subset \mathbb{R}^{pqn_\mathbf{r}}$. The penalized log-likelihood (12) can be written as

$$\ell(\boldsymbol{\psi}, \boldsymbol{\gamma}) = \ell(\boldsymbol{\psi}) - \sum_{j=1}^{p} {\boldsymbol{\psi}^j}^\mathrm{T} \mathbf{p}(\boldsymbol{\gamma}^j)\boldsymbol{\psi}^j, \tag{17}$$

where $\mathbf{p}(\boldsymbol{\gamma}^j)$ is a $n_\mathbf{r}q \times n_\mathbf{r}q$ block diagonal matrix with q blocks each of which is formed by the $n_\mathbf{r} \times n_\mathbf{r}$ matrix $\gamma_{j,k}\mathbf{s}_{j,k}$, with $\mathbf{s}_{j,k}$ a symmetric matrix s.t.

$$\int_{\mathfrak{X}_k} \{f''_{j,k}(x)\}^2 \,\mathrm{d}x = ((\mathbf{f}_k)_{j,1} \ldots (\mathbf{f}_k)_{j,n_\mathbf{r}})\mathbf{s}_{j,k}((\mathbf{f}_k)_{j,1} \ldots (\mathbf{f}_k)_{j,n_\mathbf{r}})^\mathrm{T}.$$

Note that $\mathbf{s}_{j,k}$ depends only on the observed values of the kth covariate x_{ik}, for $i = 1, \ldots, n_{\mathbf{r}}$. Based on the modified penalized log-likelihood (17), $\widehat{\boldsymbol{\psi}}$ solves the following score equations

$$\begin{pmatrix} \sum_{i=1}^{n_{\mathbf{r}}} \mathbf{m}^1\left(\mathbf{Y}_i, \boldsymbol{\psi}\right) \\ \vdots \\ \sum_{i=1}^{n_{\mathbf{r}}} \mathbf{m}^p\left(\mathbf{Y}_i, \boldsymbol{\psi}\right) \end{pmatrix} - 2 \begin{pmatrix} \mathbf{p}(\boldsymbol{\gamma}^1)\boldsymbol{\psi}^1 \\ \vdots \\ \mathbf{p}(\boldsymbol{\gamma}^p)\boldsymbol{\psi}^p \end{pmatrix} = \sum_{i=1}^{n_{\mathbf{r}}} \mathbf{m}\left(\mathbf{Y}_i, \boldsymbol{\psi}\right) - 2 \begin{pmatrix} \mathbf{p}(\boldsymbol{\gamma}^1)\boldsymbol{\psi}^1 \\ \vdots \\ \mathbf{p}(\boldsymbol{\gamma}^p)\boldsymbol{\psi}^p \end{pmatrix} = \mathbf{0}_{pqn_{\mathbf{r}}},$$

where $\mathbf{m}^j(\mathbf{y}, \boldsymbol{\psi}) = \partial \ell\left(\mathbf{y}; \boldsymbol{\psi}\right) / \partial \boldsymbol{\psi}^j$, for $j = 1, \ldots, p$. Let $\boldsymbol{\psi}_0 = (\boldsymbol{\psi}_0^{1^{\mathrm{T}}}, \ldots, \boldsymbol{\psi}_0^{p^{\mathrm{T}}})^{\mathrm{T}}$ be the true vector of parameters and $\boldsymbol{\Psi}_0$ an open neighborhood around $\boldsymbol{\psi}_0$.

Our asymptotic results hold under the following customary assumptions:

(A_1) $\boldsymbol{\gamma}^j = o(n_{\mathbf{r}}^{-1/2})\mathbf{1}_q$, for $j = 1, \ldots, p$.

(A_2) Regularity conditions:

- If $\boldsymbol{\psi} \neq \tilde{\boldsymbol{\psi}}$, then $\ell(\mathbf{y}; \boldsymbol{\psi}) \neq \ell(\mathbf{y}; \tilde{\boldsymbol{\psi}})$, with $\boldsymbol{\psi}, \tilde{\boldsymbol{\psi}} \in \boldsymbol{\Psi}$. Moreover, $\mathrm{E}\{\sup_{\boldsymbol{\psi} \in \boldsymbol{\Psi}} |\ell(\mathbf{Y}; \boldsymbol{\psi})|\} < \infty$.
- The true vector of parameters $\boldsymbol{\psi}_0$ is in the interior of $\boldsymbol{\Psi}$.
- $\ell(\mathbf{y}; \boldsymbol{\psi}) \in C^3(\boldsymbol{\Psi}_0^j)$, for $j = 1, \ldots, p$ and where $\boldsymbol{\Psi}_0 = \boldsymbol{\Psi}_0^1 \times \cdots \times \boldsymbol{\Psi}_0^p$.
- $\int \sup_{\boldsymbol{\psi} \in \boldsymbol{\Psi}_0} \|\mathbf{m}^j(\mathbf{y}, \boldsymbol{\psi})\| \, \mathrm{d}\mathbf{y} < \infty$ and $\int \sup_{\boldsymbol{\psi} \in \boldsymbol{\Psi}_0} \|\partial \mathbf{m}^j(\mathbf{y}, \boldsymbol{\psi}) / \partial \boldsymbol{\psi}^j\| \mathrm{d}\mathbf{y} < \infty$, for $j = 1, \ldots, p$.
- For $\boldsymbol{\psi} \in \boldsymbol{\Psi}_0$, $\mathrm{cov}\{\mathbf{m}^j(\mathbf{Y}, \boldsymbol{\psi})\}$ exists and is positive-definite, for all $j = 1, \ldots, p$.
- For each triplet $1 \leq q, r, s \leq n_{\mathbf{r}}$, for $j = 1, \ldots, p$, and for $\boldsymbol{\psi} \in \boldsymbol{\Psi}_0$, the quantity $|\partial^3 \ell(\mathbf{y}; \boldsymbol{\psi}) / \partial \boldsymbol{\psi}_{qrs}^j|$ is bounded from above by $M^j(\mathbf{y})$, where M^j is a real-valued function verifying $\mathrm{E}\left\{M^j(\mathbf{Y})\right\} < \infty$.

(A_3) $\mathrm{E}\left\{\mathbf{m}\left(\mathbf{Y}, \boldsymbol{\psi}\right)\right\} = \mathbf{0}_{pqn_{\mathbf{r}}}$ if and only if $\boldsymbol{\psi} = \boldsymbol{\psi}_0$.

(A_4) $\mathbf{m}^j(\mathbf{y}, \boldsymbol{\psi})$ is $C^1\left(\boldsymbol{\Psi}_0\right)$, for $j = 1, \ldots, p$.

(A_5) $\mathrm{E}\{\sup_{\boldsymbol{\psi} \in \boldsymbol{\Psi}} \|\partial \mathbf{m}^j(\mathbf{Y}, \boldsymbol{\psi}) / \partial \boldsymbol{\psi}^i\|\} < \infty$, for $i, j = 1, \ldots, p$ and $i \neq j$.

(A_6) $\mathrm{E}[\mathbf{m}^i(\mathbf{Y}, \boldsymbol{\psi}_0)\{\mathbf{m}^j(\mathbf{Y}, \boldsymbol{\psi}_0)\}^{\mathrm{T}}] = \mathbf{0}_{(qn_{\mathbf{r}}) \times (qn_{\mathbf{r}})}$, for $i, j = 1, \ldots, p$ and $i \neq j$.

The next theorem characterizes the large sample behavior of our estimator.

Theorem 1. *Under A_1–A_6, it follows that as $n_{\mathbf{r}} \to \infty$,*

1. $\|\widehat{\boldsymbol{\psi}} - \boldsymbol{\psi}_0\| = O_p(n_{\mathbf{r}}^{-1/2})$.

2. $\sqrt{n_{\mathbf{r}}}(\widehat{\boldsymbol{\psi}} - \boldsymbol{\psi}_0) \xrightarrow{\mathrm{d}} N(\mathbf{0}_{pqn_{\mathbf{r}}}, \mathbf{j}(\boldsymbol{\psi}_0)^{-1}\mathbf{v}(\boldsymbol{\psi}_0)\mathbf{j}(\boldsymbol{\psi}_0)^{-\mathrm{T}})$, *where*

$$\mathbf{j}(\boldsymbol{\psi}) = \mathrm{E}\begin{pmatrix} \frac{\partial \mathbf{m}^1}{\partial \boldsymbol{\psi}^1} & \cdots & \frac{\partial \mathbf{m}^1}{\partial \boldsymbol{\psi}^p} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{m}^p}{\partial \boldsymbol{\psi}^1} & \cdots & \frac{\partial \mathbf{m}^p}{\partial \boldsymbol{\psi}^p} \end{pmatrix} \quad \textit{and } \mathbf{v}(\boldsymbol{\psi}) = \mathrm{E}\begin{pmatrix} \mathbf{m}^1\mathbf{m}^{1^{\mathrm{T}}} & \cdots & \mathbf{m}^1\mathbf{m}^{p^{\mathrm{T}}} \\ \vdots & \ddots & \vdots \\ \mathbf{m}^p\mathbf{m}^{1^{\mathrm{T}}} & \cdots & \mathbf{m}^p\mathbf{m}^{p^{\mathrm{T}}} \end{pmatrix}.$$

3. *For* $j = 1, \ldots, p$,

$$\sqrt{n_{\mathbf{r}}}\left(\widehat{\boldsymbol{\psi}}^j - \boldsymbol{\psi}_0^j\right) \xrightarrow{d} N\left[\mathbf{0}_{qn_{\mathbf{r}}}, \left\{\mathbf{j}(\boldsymbol{\psi}_0)_{(j,j)}\right\}^{-1} \boldsymbol{\omega}_j(\boldsymbol{\psi}_0) \left\{\mathbf{j}(\boldsymbol{\psi}_0)_{(j,j)}\right\}^{-\mathrm{T}}\right],$$

where

$$\boldsymbol{\omega}_j(\boldsymbol{\psi}) = \mathrm{cov}\left[\mathbf{m}^j - \left(\frac{\partial \mathbf{m}^j}{\partial \boldsymbol{\psi}^1} \cdots \frac{\partial \mathbf{m}^j}{\partial \boldsymbol{\psi}^{j-1}} \frac{\partial \mathbf{m}^j}{\partial \boldsymbol{\psi}^{j+1}} \cdots \frac{\partial \mathbf{m}^j}{\partial \boldsymbol{\psi}^p}\right) \left\{\mathbf{j}(\boldsymbol{\psi})_{(-j,-j)}\right\}^{-1} \left(\mathbf{m}^1 \cdots \mathbf{m}^{j-1}\mathbf{m}^{j+1} \cdots \mathbf{m}^p\right)^{\mathrm{T}}\right],$$

$\mathbf{j}(\boldsymbol{\psi})_{(j,j)}$ *is the* j*th* $qn_{\mathbf{r}} \times qn_{\mathbf{r}}$ *block matrix of* $\mathbf{j}(\boldsymbol{\psi})$, *and* $\mathbf{j}(\boldsymbol{\psi})_{(-j,-j)}$ *is the* $\{(p-1)qn_{\mathbf{r}}\} \times \{(p-1)qn_{\mathbf{r}}\}$ *matrix constructed from* $\mathbf{j}(\boldsymbol{\psi})$ *by dropping the* j*th* $qn_{\mathbf{r}} \times qn_{\mathbf{r}}$ *block matrix.*

These results are derived from a Taylor expansion of the penalized log-likelihood along the same lines as in Vatter and Chavez-Demoulin (2015) and Davison (2003, p. 147). Similar results on the large sample behavior of the corresponding plug-in estimator (16) can be derived using the multivariate delta method. These results are useful to derive and construct approximate confidence intervals for conditional angular densities, and compare nested models based on likelihood ratio tests.

Furthermore, it should be noted that our proviso is similar to that of de Carvalho and Davison (2014) in the sense that asymptotic properties of the estimator $\widehat{\boldsymbol{\psi}}$ are derived under the assumption of known margins and that we sample from the limiting object $h_{\mathbf{x}}$ whereas in practice only a sample of (estimated) pseudo-angles, $\{\widehat{W}_i\}_{i=1}^{n_{\mathbf{r}}}$, would be available. Asymptotic properties under misspecification of the parametric model set for $h_{\mathbf{x}}$ could in principle be derived under additional assumptions on $\boldsymbol{\psi}$ and $\mathbf{m}$, along the same lines as in standard likelihood theory (Knight, 2000). The resulting theory is outside the scope of this work and is deliberately not studied here.

4 SIMULATION STUDY

4.1 Monte Carlo Evidence in the Bivariate Setting

We start by assessing the performance of our covariate-dependent spectral density estimator using the bivariate extremal dependence structures presented in Section 2.2. For each dependence model

from Examples 1–3, we draw 500 samples $\{(Y_{i,1}, Y_{i,2})\}_{i=1}^{n}$ with sample size $n = 6000$ and where each observation $(Y_{i,1}, Y_{i,2})$ has unit Fréchet margins and is drawn from the chosen dependence model conditional on a fixed value x_i of a covariate x.

The parameters of h_x being restricted to be in $[0, 1]$ in Example 1 and in $(0, \infty)$ in Examples 2, and 3, there is a need for the use of a link function when assessing the dependence of these parameters on the covariate x. We use the inverse logit link function whenever the parameter is restricted to be in $[0, 1]$ and the exponential function whenever the parameter needs to be strictly positive.
For estimating h_x, we only consider the observations whose radial component exceeds its 95% quantile and we hence end-up with $n_{\mathbf{r}} = 300$ extreme observations. To gain insight into the bias and variance of our covariate-adjusted spectral density estimator, randomly selected bootstrap estimates along with the Monte Carlo means are considered. The resulting bias and variance are only induced by our estimation technique and reflect neither the uncertainty due to the asymptotic assumption on the distribution of the angular observations nor the choice of the parametric model.

Figure 2 displays the trajectories of 100 randomly selected estimates of the covariate-adjusted spectral densities from Examples 1, 2, and 3 for various fixed values of the covariate x, along with the Monte Carlo means based on the 500 bootstrap samples. All panels show that the Monte Carlo means produce relatively accurate estimations of the spectral density. A slight upward bias in the estimates is however observed when approaching extremal independence.

4.2 Monte Carlo Evidence in the Three-Variate Setting

We now consider the case of a trivariate covariate-dependent logistic dependence structure. To guarantee a valid logistic spectral density estimate, we consider the same model as in (10) but in the three dimensional case, i.e.,

$$h_x(\mathbf{w}) = \prod_{i=1}^{2}(i/\alpha_x - 1)\left(\prod_{i=1}^{3} w_i\right)^{-1-1/\alpha_x}\left(\sum_{i=1}^{3} w_i^{-1/\alpha_x}\right)^{\alpha_x - 3}, \quad \mathbf{w} = (w_1, w_2, w_3) \in S_3, \tag{18}$$

where $\alpha_x = x$, with $x \in [0.2, 0.9]$. We draw 500 bootstrap samples $\{(Y_{i,1}, Y_{i,2}, Y_{i,3})\}_{i=1}^{n}$ with sample size $n = 6000$ and where each observation $(Y_{i,1}, Y_{i,2}, Y_{i,3})$ has unit Fréchet margins and is drawn from a symmetric logistic model conditional on a fixed value x_i of the covariate x. As in Example 1, the inverse logit link function is imposed in (18) to ensure the right range of values of the parameter α_x. We again threshold the radial component of the simulated observations and retain the observations with a radial component exceeding its 95% quantile. Figure 3 displays the contour plots of the Monte Carlo mean of the covariate-adjusted spectral density in (18) at three fixed values of x so

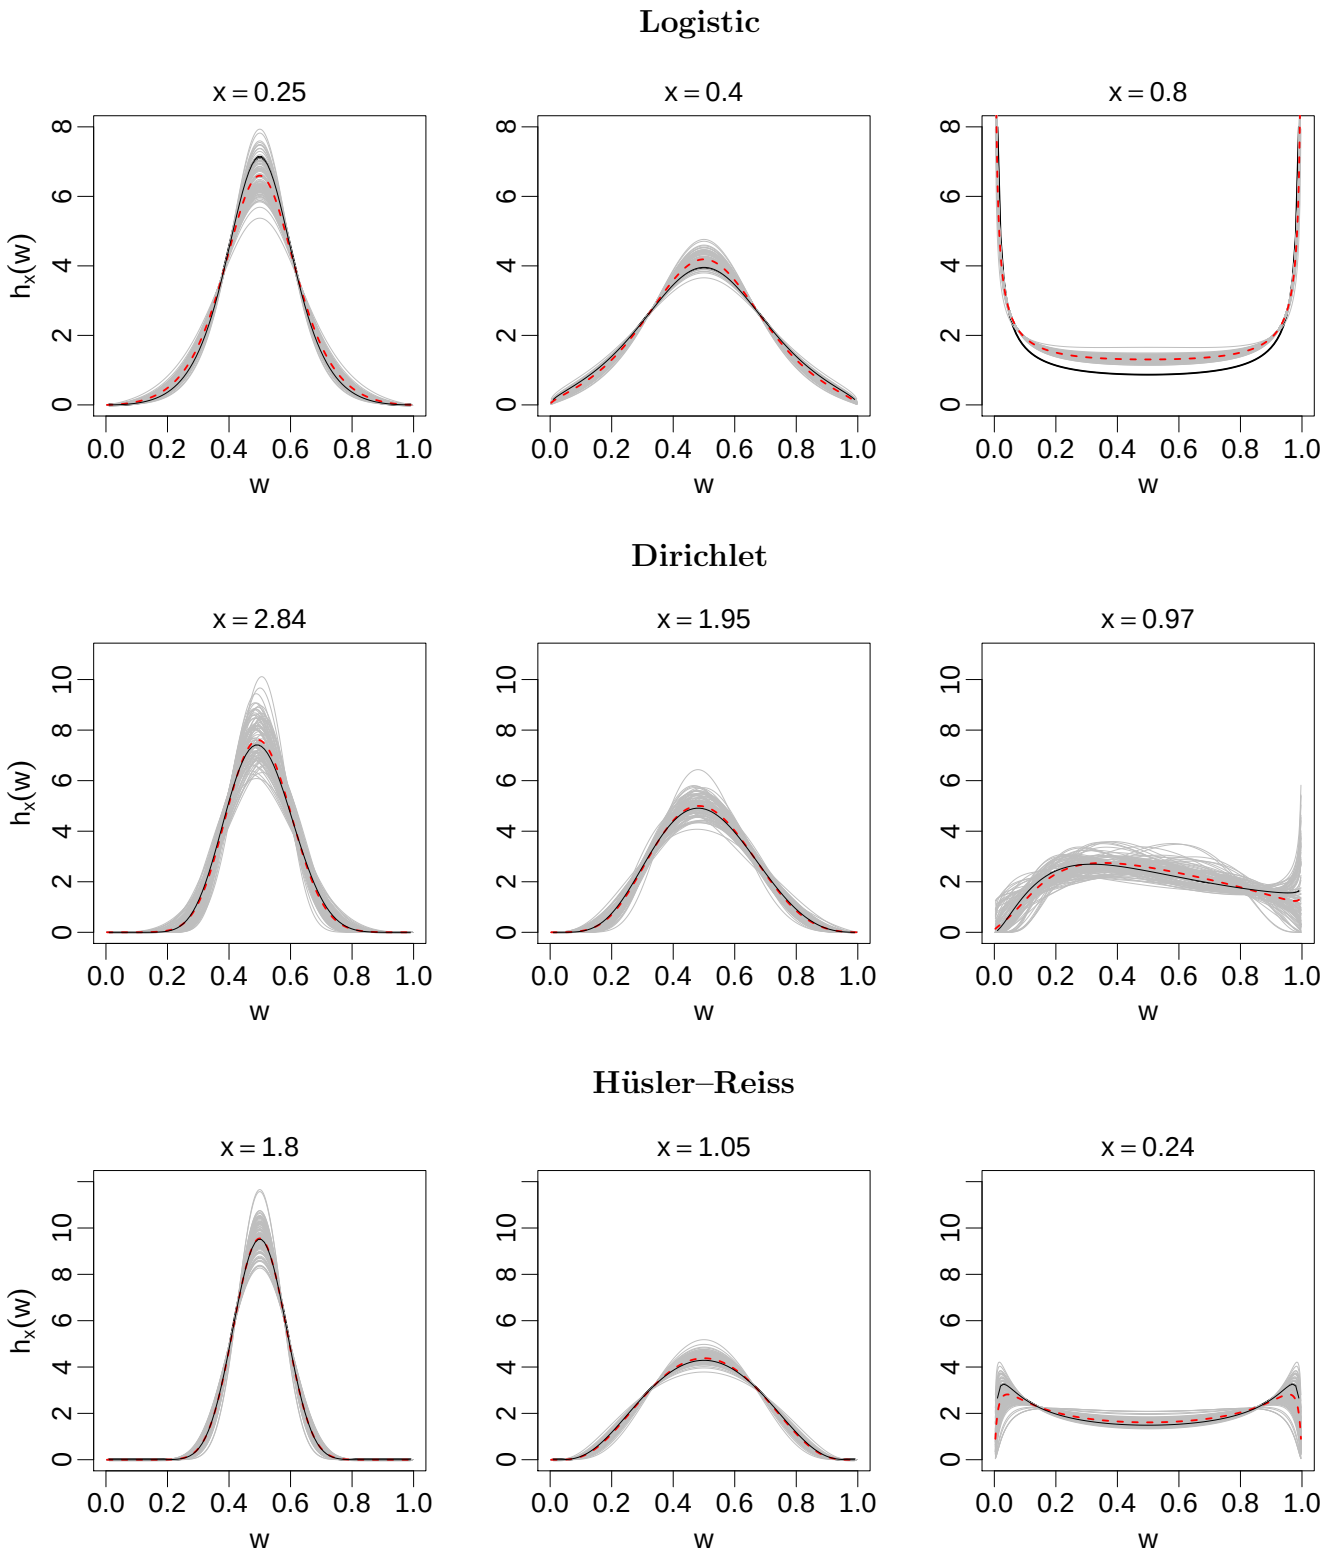

Figure 2: Randomly selected covariate-adjusted estimates of the spectral densities in Examples 1, 2, and 3 conditional on different values of the covariate x (in grey) along with their Monte Carlo means (dashed red lines). The true spectral densities are displayed in solid lines.

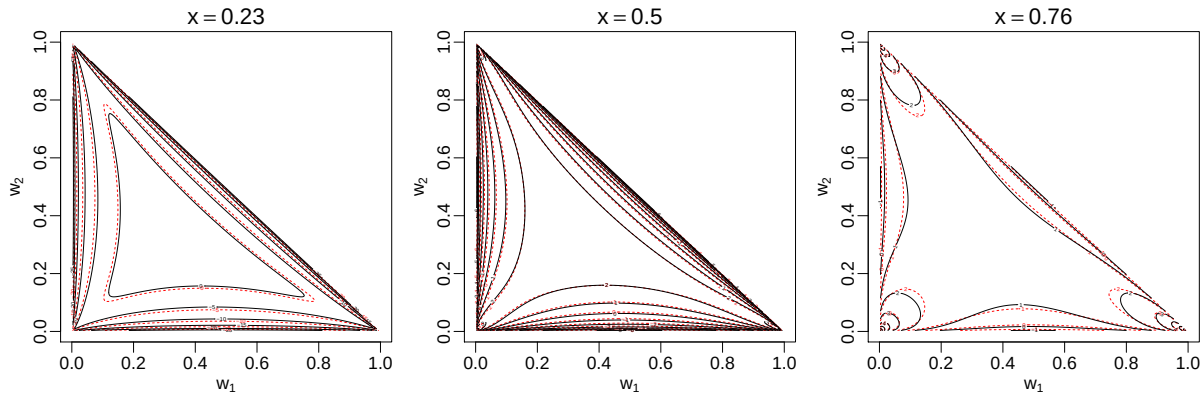

Figure 3: Contour plots of the Monte Carlo mean of the three-variate covariate-adjusted spectral density in (18) (dashed red lines) at three fixed values of x. The contour plots of the true spectral density are displayed in solid black lines.

that different extremal dependence schemes are treated. All panels in Figure 3 show that, for the different extremal dependence strengths, the contour plots of the Monte Carlo mean spectral density are remarkably close to the actual contour plots. As for the bivariate case, a small bias is observed when approaching extremal independence.

5 TRACKING EXTREMAL DEPENDENCE

5.1 Data Description and Motivation for the Analysis

Today's global financial landscape calls for the need to develop more accurate inferences on the likelihood of certain risks. Figure 4 shows joint occurrences of extreme losses for two major players in the banking industry—Citibank and JP Morgan—during two different time periods as well as the joint extreme losses for two key players in the healthcare industry—Merk & Co and Patterson Companies.

A comparison of the left and right plots in Figure 4 suggests that, under the assumption that the two banking stock markets are asymptotically dependent, the strength of the extremal dependence is different before and during the subprime crisis. Additionally, a comparison of the top and bottom plots suggests that, during the same time period, the joint extreme losses in the banking sector behave differently than those in the healthcare sector. Hence, there is a need for extremal dependence models that account for different extremal dependence structures over time and between sectors.

We apply our methodology to estimate the dynamics between extreme events for financial institutions. We consider two sectors and four stocks within each of these sectors: the banking

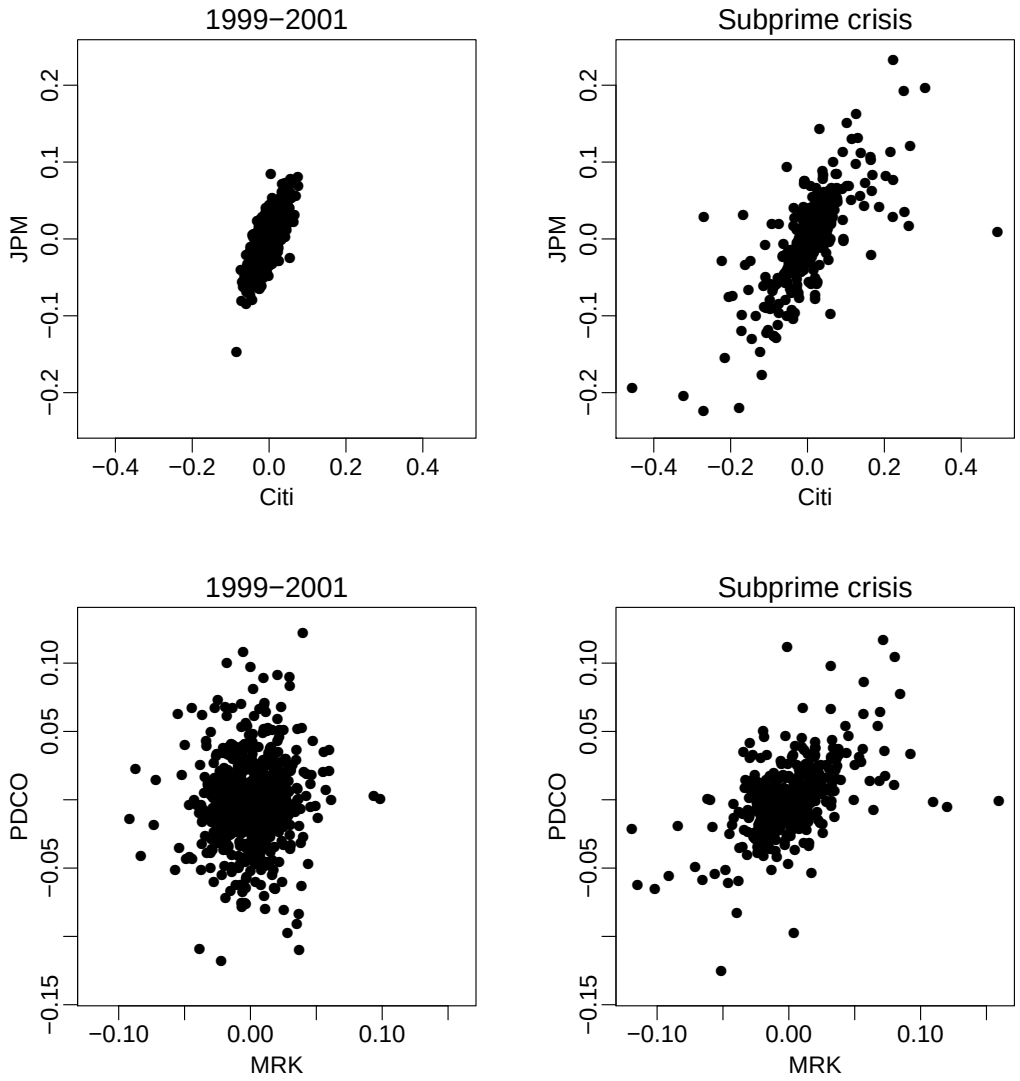

Figure 4: Scatterplots of negative daily log-returns stock prices of Citibank (CITI), JP Morgan (JPM), Merck & Co (MRK), and Patterson Companies (PDCO). Left: Time period of 1999–2001; Right: Time period corresponding to the subprime crisis.

sector encompassing the Citi Group (Citi), Comerica (CMA), the Fifth Third Bancorp (FITB), and J.P. Morgan (JPM), along with the healthcare sector encompassing Biogen (BIIB), Express Scripts Holding Company (ESRX), Merck & Co (MRK), and Patterson Companies (PDCO). The motivation for a simultaneous analysis of these two sectors is that they are very different in terms of business cycle fluctuations. On one hand, when the business cycle is in an upswing phase, the banking sector benefits from additional investments due to improved economic conditions and an increase in personal and capital investing. On the other hand, with a consistent demand for goods and services, the healthcare sector is less sensitive to business cycle fluctuations. The dataset consists of the negative log-returns on the daily closing prices from January 6th, 1998 to December 31st, 2015 and results in $n = 4523$ observations for each stock. The data were obtained from Yahoo Finance (`https://uk.finance.yahoo.com`). The loss returns recorded during two non-overlapping time periods for four stocks in the different sectors are displayed in Figure 4. The advantage of our method over the usual estimation techniques applied to extremal dependence is that we borrow information from the different sectors and different time periods by pooling the data and then estimating the extremal dependence as a function of various covariates in the different sectors and over time. Pooling together the data increases the amount of information and provides significance relative to all the different types of data included.

5.2 Marginal Modeling of the Stocks

Prior to the estimation of the extremal dependence structure, we transform the margins to a common unit Fréchet scale. A stylized feature of financial time series needs to be taken into account during the marginal modeling procedure. In fact, daily returns of assets or portfolios are known to be conditionally heteroskedastic (Lamoureux and Lastrapes, 1990) and hence need to be modeled using the ARCH (Engle, 1982) and GARCH (Bollerslev, 1986) models which properly take into account this clustering of volatility, see (Tsay, 2005, ch. 3) for details. For each individual stock i, for $i = 1, \ldots, 8$, we fit the widely-used GARCH$(1, 1)$ (McNeil et al., 2015, sec. 4.2) model to the negative log-returns $Y_i = (Y_{i,1}, \ldots, Y_{i,n})$:

$$\begin{aligned} Y_{i,j} &= \mu + \sigma_{i,j}\varepsilon_{i,j} \\ \sigma^2_{i,j} &= \alpha_{i,0} + \alpha_{i,1}\left(Y_{i,j-1} - \mu_i\right)^2 + \beta_{i,1}\sigma^2_{i,j-1}, \quad j = 2, \ldots, n, \end{aligned} \tag{19}$$

where $\boldsymbol{\varepsilon}_i = \{\varepsilon_{i,j}\}_{j=1}^n$ is a white noise with zero mean and unit variance. A proper model selection based on AIC or BIC could be performed on the series of negative log-returns for each stock to choose the right order of the GARCH model. However, it is often argued in the literature that

the parsimonious GARCH$(1,1)$ model is sufficient in most applications where the volatility is to be removed (Hansen and Lunde, 2005).

We estimate our model in (19) using quasi maximum likelihood (see, e.g, Elie and Jeantheau, 1995). This estimation procedure allows for a distribution-free model for the innovations $\boldsymbol{\varepsilon}_i$ hence avoiding any inconsistency that would result from the misspecification of the innovations' distribution. Diagnostic plots such as partial auto-correlation function plots are used to assess the validity of the fitted model. The resulting series of residuals for each asset are extracted and transformed marginally to unit Fréchet using the following semi-parametric procedure: For each individual stock i, $i = 1, \ldots, 8$, an arbitrary threshold u_i equal to the 98% empirical quantile of $\boldsymbol{\varepsilon}_i$ is fixed and the marginal distribution function is estimated by $\widehat{F}_i$ which is equal to a generalized Pareto distribution (GPD) fitted to the data above u_i and to the empirical distribution function below u_i.

A graphical goodness-of-fit test for the GPD model fitted to the threshold excesses consists in comparing the distribution of a test statistic S with the unit exponential distribution: if a random variable Y follows a GPD$(\sigma,\, \xi)$ denoted $G_{\sigma,\xi}$, then $S = -\ln\{1 - G_{\sigma,\xi}(Y)\}$ is unit exponentially distributed. Figure 5 displays the resulting qq-plots for the following four assets: CMA, JPM, MRK, and PDCO. The marginal modeling of the threshold excesses looks fine and we now proceed to the transformation of the series of residuals to a common unit Fréchet scale by applying the transformation $Z_{i,j} = -[\ln\{\widehat{F}_i(\varepsilon_{i,j})\}]^{-1}$, for $j = 1, \ldots, n$ and $i = 1, \ldots, 8$.

Parenthetically, we note that the validity of the above semi-parametric procedure relies on the stationarity of the series of residuals resulting from the GARCH modeling procedure. Indeed, if the negative log-returns were to be transformed to unit Fréchet without modeling and removing the clustering of volatility, one would have had to model the bulk of the data, i.e., the observations below the fixed threshold, using an arbitrary chosen distribution with time-dependent parameters, and the threshold excesses using the generalized Pareto distribution with a GAM fitted to its parameters (see Chavez-Demoulin and Davison, 2005, for details). However, this technique can produce misleading results due to the arbitrary choice of the distribution fitted to the bulk of the data.

5.3 Presence of Extremal Dependence

To verify that an application of our methods does not lead to an overestimation of risk, we need to assess whether the marginal residuals are asymptotically dependent. To do so, two well-known measures can be computed: the χ and $\bar{\chi}$ dependence measures (Coles et al., 1999; Heffernan, 2000). The coefficient of extremal dependence χ reflects the strength of dependence at extreme levels, provided the data are asymptotically dependent. The dependence measure $\bar{\chi}$ reflects the

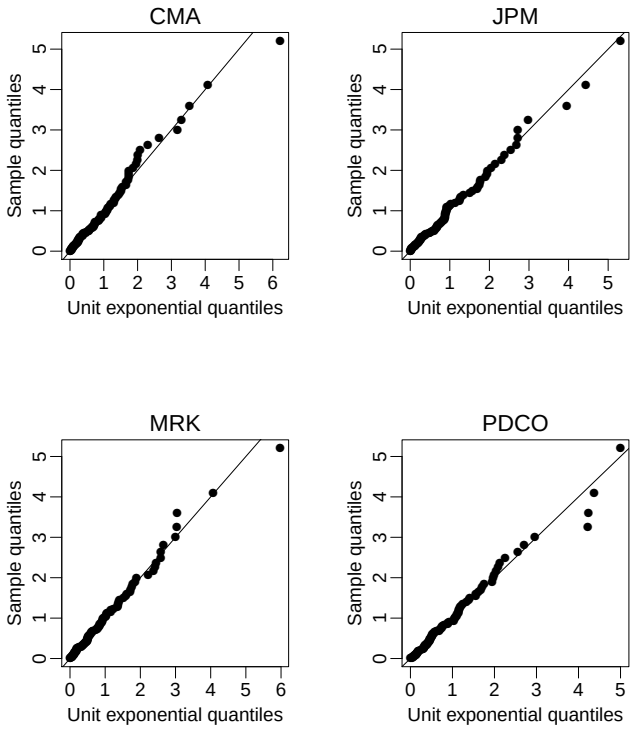

Figure 5: Diagnostic plots of the GPD modeling of the residuals' threshold excesses for the CMA and JPM stocks in the banking sector and the MRK and PDCO stocks in the healthcare sector.

strength of dependence when the data are asymptotically independent. If $\bar{\chi} = 1$, the data are asymptotically dependent and the corresponding value of χ measures the strength of extremal dependence. However, if $\bar{\chi} < 1$, the data are asymptotically independent and the corresponding value of χ is of no interest in the analysis of extremal dependence. Empirical estimates of the χ and $\bar{\chi}$ measures can be computed and result in two diagnostic plots: the χ-plot and the $\bar{\chi}$-plot, respectively.

We consider the series of residuals $\boldsymbol{\varepsilon}_i$ of a pair of stocks in each of the banking and healthcare sectors and compare the empirical estimates of their dependence measures χ and $\bar{\chi}$ taken over different time periods. In the banking sector, it is interesting to investigate the extremal dependence before, during, and after the subprime crisis which we attribute to the period from December 1st, 2007 to June 30th, 2009. In the healthcare sector, we choose two arbitrary time periods, in the beginning and in the end of the period of analysis. Figures 6 and 7 display the χ-plots and the $\bar{\chi}$-plots over different time periods of the pairs Comerica–J.P. Morgan (CMA–JPM) and Merck–Co–Patterson Companies (MRK–PDCO), respectively.

The interpretation of the $\bar{\chi}$-plots of the pair CMA–JPM is not straightforward. However, we

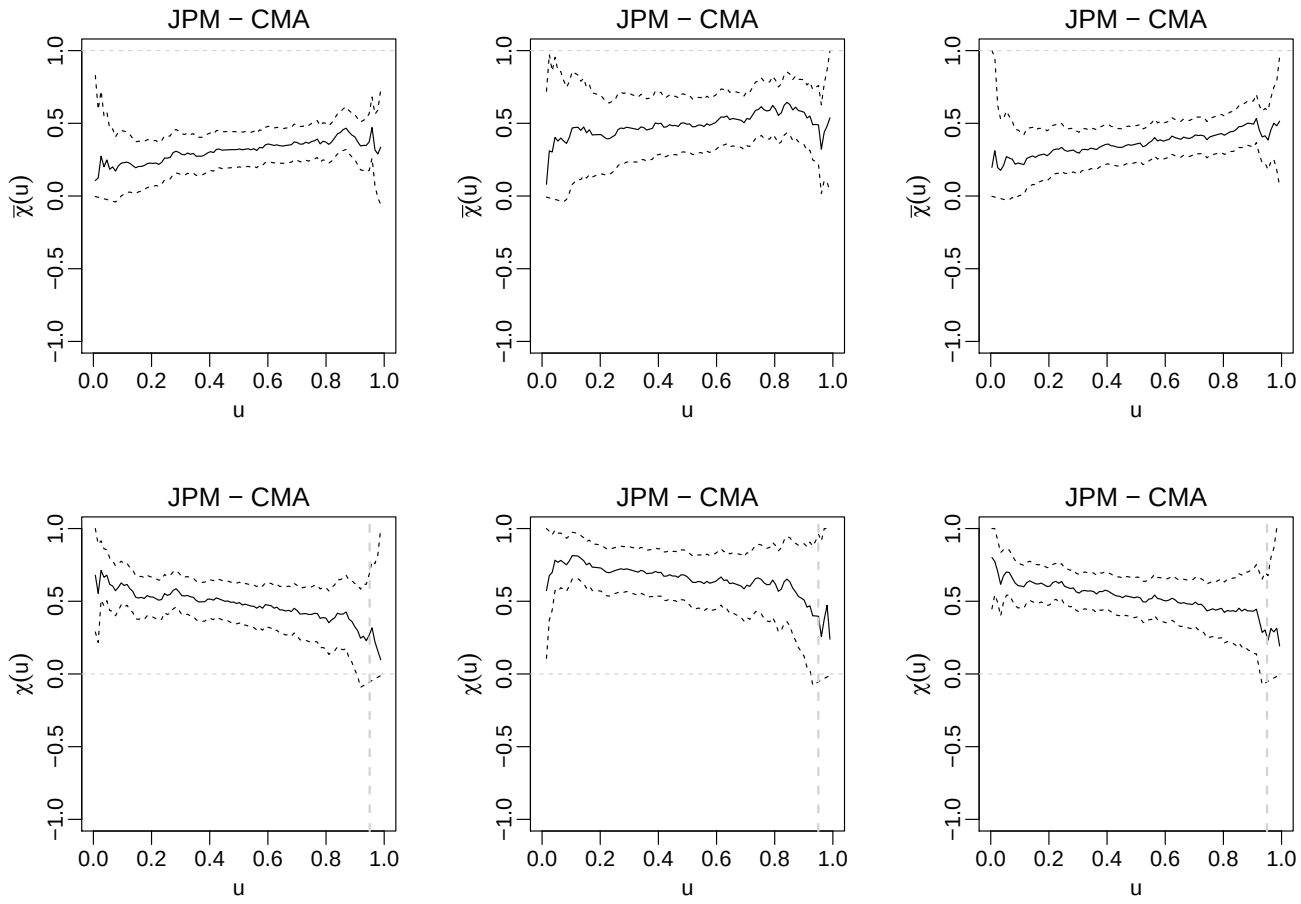

Figure 6: $\bar{\chi}$-plots (top panels) and χ-plots (bottom panels) of the pair CMA–JPM during the periods 1999–2001 (left panels), the subprime crisis (middle panels), and 2012–2014 (right panels).

observe that in the $\bar{\chi}$-plot in the first time period, the 95% confidence intervals do not contain the value 1 as opposed to the $\bar{\chi}$-plots in the second and third time periods, where $\bar{\chi}(u)$ seems to tend to 1 as $u \to 1$. Regarding the χ-plots in the second and third time periods, when comparing the values of $\chi(0.95)$ (indicated by the dashed grey lines) in both plots, it appears that the strengths of extremal dependence are different with a value around 0.4 during the subprime crisis and a value around 0.2 after the crisis. Unsurprisingly, extreme losses of these stocks were more dependent during the crisis than after the crisis when financial regulatory reforms were addressed (Claessens and Kodres, 2014). Figure 7 can be interpreted in the same way. The $\bar{\chi}$-plots reflect an extremal independence between the pair of stocks MRK–PDCO during the first time period, while the marginal residuals of these two stocks seem to be asymptotically dependent during the second period with a value of $\chi(0.95)$ around 0.3.

Hence, we proceed by assuming the presence of extremal dependence in the data.

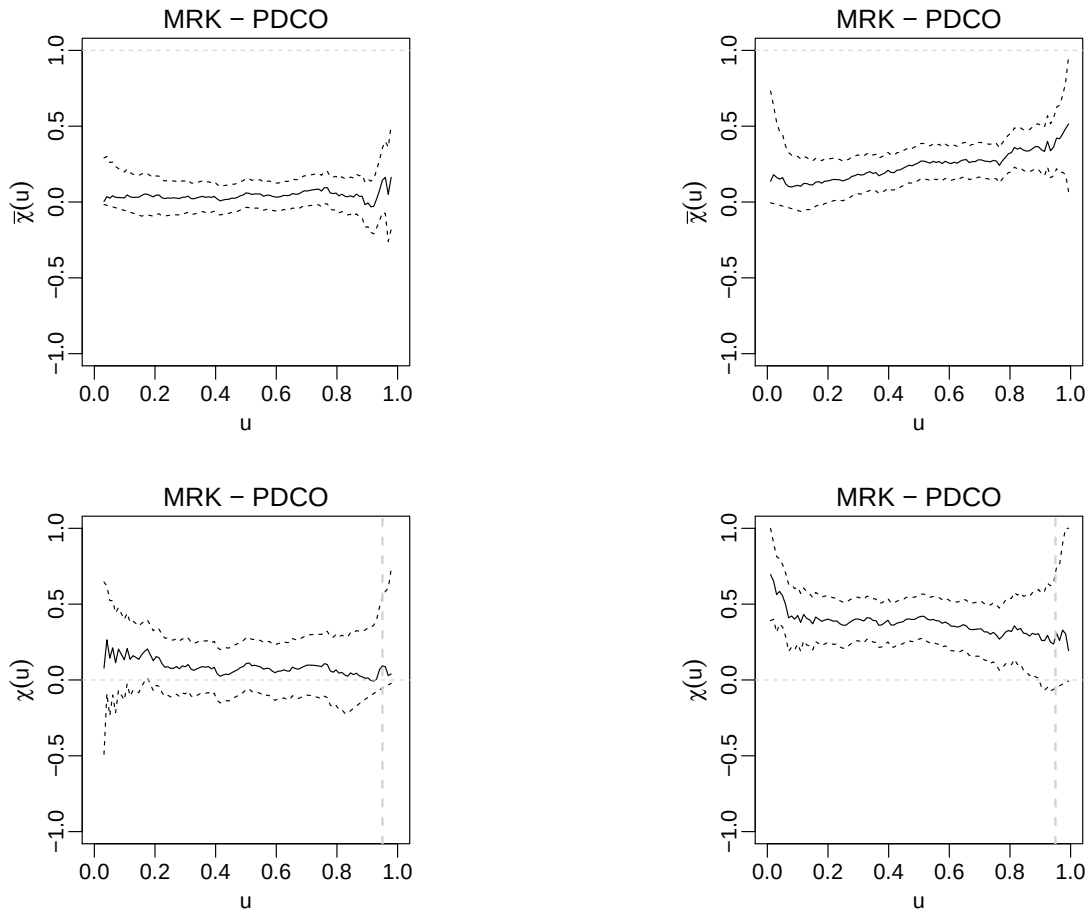

Figure 7: $\bar{\chi}$-plots (top panels) and χ-plots (bottom panels) of the pair MRK–PDCO during the time periods 1999–2001 (left panels) and 2011–2013 (right panels).

5.4 Cross-Sectoral and Time-Varying Extremal Dependence

We now focus on the estimation of the extremal dependence between the different assets, as a function of the sector's type and time. The motivation behind this choice of covariates is twofold: as depicted from the χ-plots in Figure 6, the dependence between extreme events is likely to vary over time due to events such as financial crises, governmental policy reforms, etc. Second, as stated above, extreme events for each sector are likely to behave differently and most of all to depend on different causes. During the observed period of time the two main causes which affected the two sectors were the financial crisis and the Obamacare reform. The financial crisis affected the banking sector over the short term, by causing banks to lose large amounts of money on mortgage defaults. As a consequence new regulatory actions were put in place through Basel III. Obamacare affected the healthcare sector in a different way: quoting CNBC[††] on June, 2 2015:

[††] http://www.cnbc.com/2015/06/02/health-care-stocks-a-great-long-term-market-bet.html

Table 1: Maximum likelihood estimates of the logistic dependence parameter with asymptotic standard errors in parentheses.

	1999–2001	Subprime crisis
Citi vs JPM	$0.73_{(0.64)}$	$0.68_{(0.61)}$
MRK vs PDCO	$0.81_{(0.51)}$	$0.71_{(0.57)}$

> "In 2010, the health-care industry was in turmoil. President Barack Obama had just signed the Affordable Care Act into law, and there was significant uncertainty around how this would impact the industry. That year, health care wasn't just the worst-performing sector—it was the only one to post negative returns... Slowly but surely, though, it became clear that the Affordable Care Act was going to boost company bottom lines, not hurt them..."

and, in the same article it is written:

> "The health-care sector dominance is continuing, with year-to-date returns of 8.8 percent, the best performance of any sector, and there's good historical reason to believe the run won't be ending."

An informal way to validate this choice of covariates is to model the extremal dependence between the assets represented in Figure 4 in the different sectors and in the two considered time periods. Due to its simplicity of implementation and ease of interpretability, we fit the symmetric logistic model to the extreme angular observations and compare the resulting 'stationary' dependence parameters. Table 1 gives the maximum likelihood estimators of the logistic dependence parameters along with their standard errors.

As pointed out previously, based on the plots in Figure 4, the dependence parameter estimates vary through time and are different from one sector to another. The MLEs are based on a small amount of data in each case (26 extreme observations in the first time period and 20 during the subprime crisis) which results in very large standard errors. This issue can be dealt with by pooling the data and modeling a covariate-dependent logistic parameter.

The analysis is now conducted on the pooled data with a logistic time-dependent parameter that is different depending on whether the observed pair of stocks lies in the banking sector or the healthcare sector. Again, as in Section 4, the inverse logit link is imposed on the dependence parameter to ensure a valid angular density. Even though the logistic model is known to be rigid, the generalized additive modeling of its parameter, adds some flexibility while maintaining the desired simplicity.

In each sector, we retain the two-dimensional observations whose radial component is large in the sense that it exceeds a large empirical threshold fixed at the 95% quantile. One can argue that

since the (asymptotic) distribution of the angular observations is to be modeled as non-stationary over time, a time-dependence of the radial component in each sector should be taken into account. A parametric quantile regression at the 95% level is performed on the two radial components and the effect of time on extreme radial observations is found to be not significant in both sectors. Figure 8 shows the dynamics of the logistic dependence parameter through time and in each sector. Point estimates along with Monte Carlo means (based on 1000 bootstrap samples) are depicted. Pointwise asymptotic 95% confidence intervals, computed based on a Monte Carlo integration, are also shown and compared to the smoothed bootstrap confidence intervals. Both confidence intervals reflect the uncertainty resulting from the estimation of the parameter of the angular density and account for neither the uncertainty due to the marginal estimation nor the choice of the ultimate parametric model.

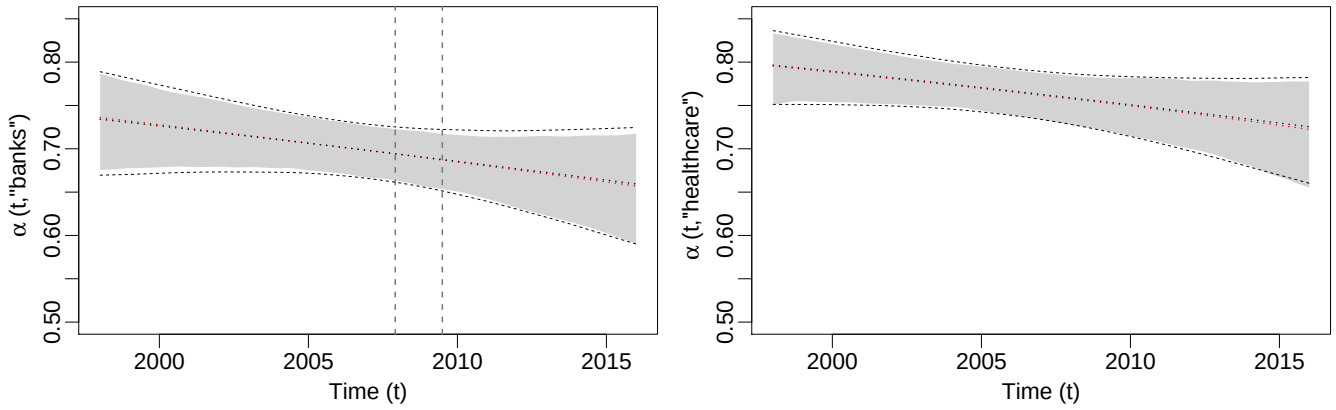

Figure 8: GAM for the bivariate logistic dependence parameter as a smooth function of time in the banking sector (left panel) and the healthcare sector (right panel). The vertical dashed grey lines correspond to the subprime crisis. Point estimates (dotted black lines), Monte Carlo means (dotted red lines), and bootstrap (grey areas) and asymptotic (dashed black lines) 95% confidence intervals are represented.

The estimates are consistent with the MLEs in Table 1. As expected, the asymptotic uncertainty is smaller than in the 'stationary' case due to the augmented number of considered observations (227 extreme observations in each sector).

The entire dataset consisting of the eight assets described in Section 5.1 is now considered. In each sector, we retain the four-dimensional observations with a radial component exceeding its 95% quantile. The resulting extreme observations are displayed in Figure 9 according to the type of sector. The consequences of the subprime crisis (from December 1st, 2007 until June 30th, 2009) can easily be observed in the banking sector where the negative log-returns of the four stocks are

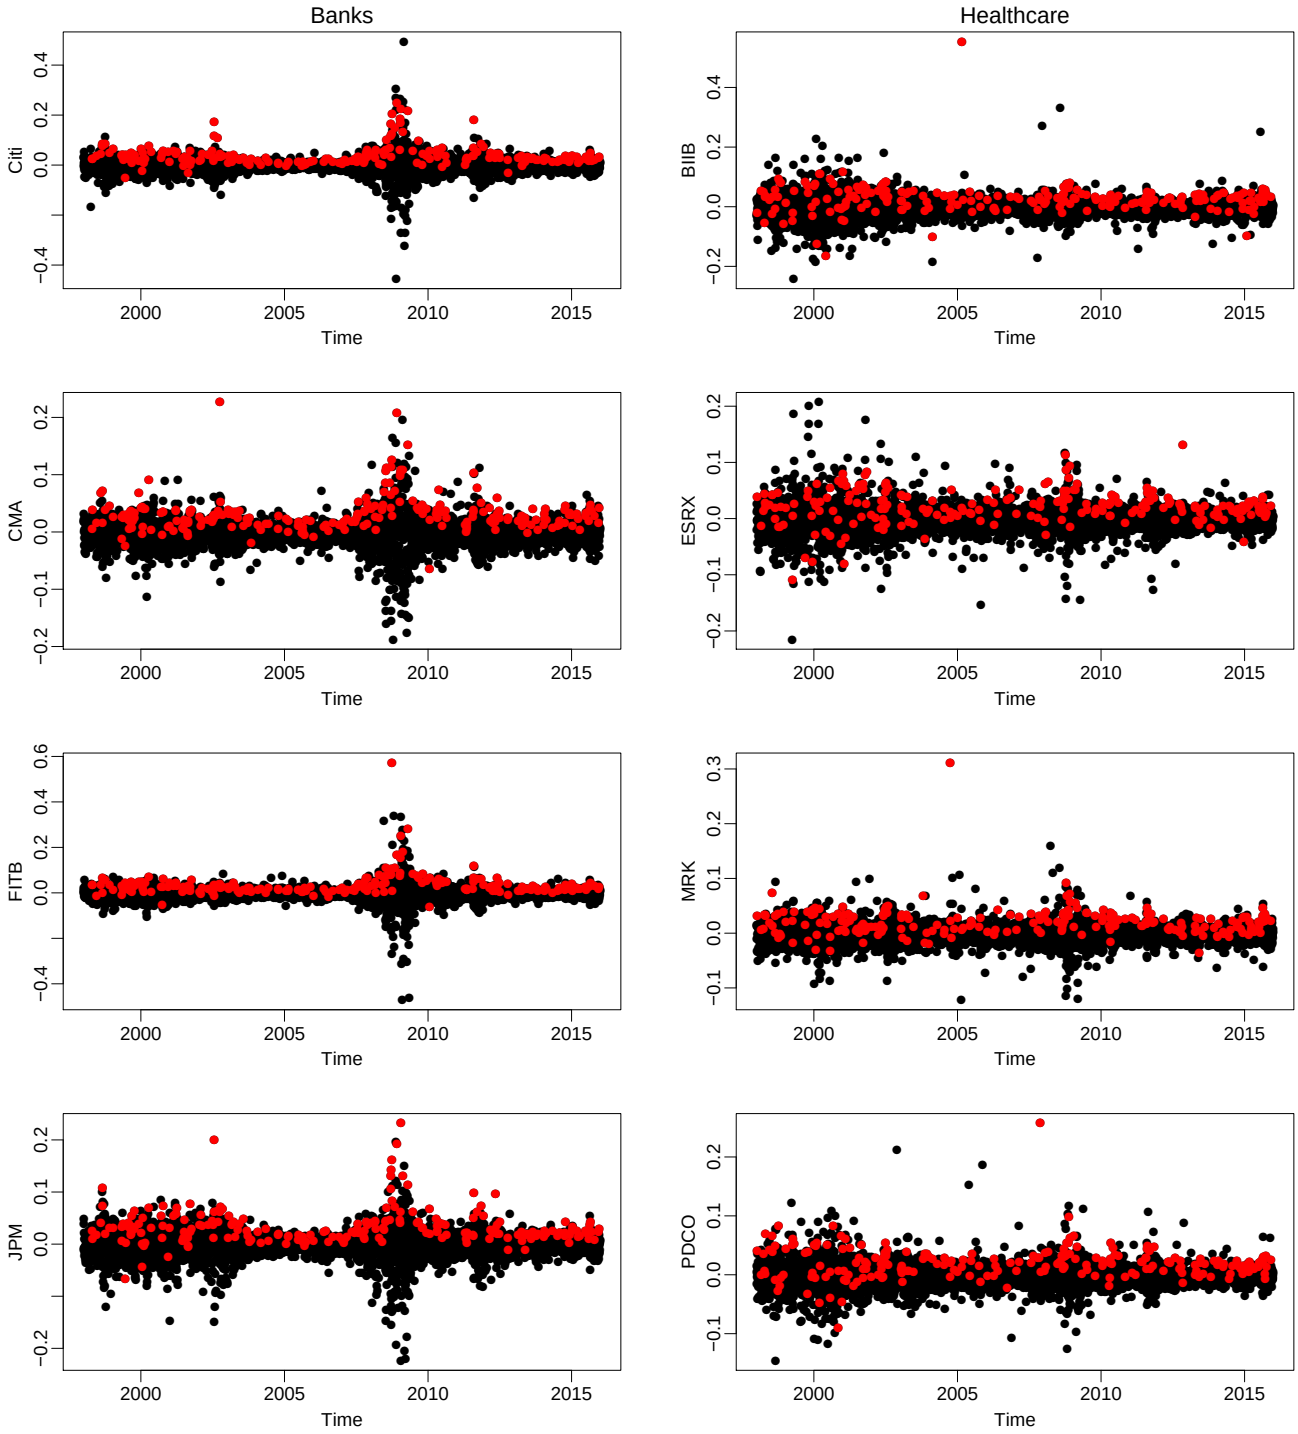

Figure 9: Stocks' negative log-returns and the retained extreme observations (red filled dots).

unusually high. We can also observe a stock market downturn in 2002 which only affected Citi and JPM in July and CMA in October, while FITB was unaffected. Finally, compared to the banking sector, the stocks in the healthcare sector seem to react differently to the subprime crisis which had less impact and they appear to be affected by different stimuli. The four-dimensional covariate-adjusted logistic model is fitted to the extreme observations and its dependence parameter is displayed in Figure 10. The dynamics of the dependence between the extreme losses in the

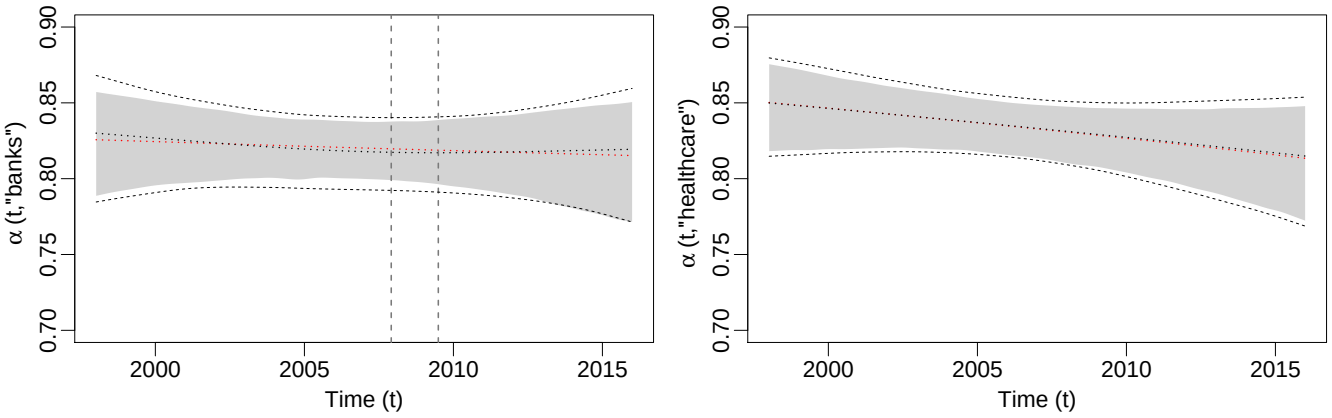

Figure 10: GAM for the four-dimensional logistic dependence parameter as a smooth function of time in the banking sector (left panel) and the healthcare sector (right panel). The vertical dashed grey lines correspond to the subprime crisis. Point estimates (dotted black lines), Monte Carlo means (dotted red lines), and bootstrap (grey areas) and asymptotic (dashed black lines) 95% confidence intervals are represented.

banking sector are characterised by an increase in the extremal dependence during the subprime financial crisis and a large uncertainty following the crisis. Even though banking regulations were implemented after the financial crisis, the extremal dependence seems to have stabilized around the minimum value recorded on July 7, 2009. The dynamics of the dependence between the extreme losses in the healthcare sector are however different. A slight but significant negative trend is observed during the studied time period. The increase in the extremal dependence is higher in the healthcare sector than in the banking sector. This increase might be explained by the fact that the entire healthcare sector is heavily affected by government decisions compared to other sectors. In fact, Koijen et al. (2016) showed that large losses in the healthcare sector are observed when there are severe threats of government intervention.

Finally, it should be noted that applying the inverse logit link function on the logistic covariate-dependent parameter imposes a high degree of smoothness on its shape, which filters out the short-term variations in the extremal dependence.

6 FINAL REMARKS

In this paper, we have introduced a sturdy and general approach to model the influence of covariates on the extremal dependence structure. Keeping in mind that extreme values are scarce, our methodology borrows strength from a parametric assumption and benefits directly from the flexibility of vector generalized additive models. An important advantage over existing methods is that our model profits from the resilience of the GAM paradigm allowing the incorporation of a large number of covariates of different types (continuous, factor, etc). The fitting procedure is an iterative ridge regression, the implementation of which is based on a Newton–Raphson type algorithm that is available in many statistical software. An illustration is provided in the R code in the Supplementary Material.

The method paves the way for novel applications, in line with topical questions and concerns in a wide variety of fields. The approach is valid in high-dimensions. However, the number of parameters would increase quickly with the dimension and additional complications would arise. Relying on composite likelihoods (Padoan et al., 2010) instead of the full likelihood seems to represent a promising path for future extensions of the proposed methodology in a high-dimensional context.

SUPPLEMENTARY MATERIAL

The online supplement to this article contains an R code for implementing the penalized maximum likelihood estimation of the dependence parameter of the covariate-adjusted Hüsler–Reiss angular density.

Fitting of the covariate-adjusted Hüsler–Reiss angular density: Code to perform the Newton–Raphson algorithm needed for the estimation of the covariate-adjusted Hüsler–Reiss angular density. The estimation procedure is illustrated by reproducing the results of the simulation study under the setting of Example 3. (Weaving Rnw file)

References

Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J., De Waal, D. and Ferro, C. (2004), *Statistics of Extremes: Theory and Applications*, Wiley, New York.

Boldi, M.-O. and Davison, A. C. (2007), 'A mixture model for multivariate extremes', *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **69**(2), 217–229.

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **31**, 307–327.

Castro, D. and de Carvalho, M. (2017), 'Spectral density regression for bivariate extremes', *Stochastic Environmental Research and Risk Assessment* .

Chavez-Demoulin, V. and Davison, A. C. (2005), 'Generalized additive modelling of sample extremes', *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **54**(1), 207–222.

Chavez-Demoulin, V., Embrechts, P. and Hofert, M. (2015), 'An extreme value approach for modeling operational risk losses depending on covariates', *Journal of Risk and Insurance* **83**, 735–776.

Claessens, S. and Kodres, L. (2014), 'The regulatory responses to the global financial crisis: Some uncomfortable questions', *IMF Working Papers* .

Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, Springer, London.

Coles, S., Heffernan, J. E. and Tawn, J. A. (1999), 'Dependence measures for extreme value analyses', *Extremes* **2**, 339–365.

Coles, S. and Tawn, J. A. (1991), 'Modelling Extreme Multivariate Events', *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **53**, 377–392.

Davison, A. C. (2003), *Statistical Models*, Cambridge University Press, Cambridge, UK.

de Carvalho, M. (2016), Statistics of Extremes: Challenges and Opportunities, *in* F. Longin, ed., 'Extreme Events in Finance: A Handbook of Extreme Value Theory and Its Applications', Wiley, Hoboken.

de Carvalho, M. and Davison, A. C. (2014), 'Spectral Density Ratio Models for Multivariate Extremes', *Journal of the American Statistical Association* **109**(506), 764–776.

de Carvalho, M., Oumow, B., Segers, J. and Warchoł, M. (2013), 'A Euclidean Likelihood Estimator for Bivariate Tail Dependence', *Communications in Statistics—Theory and Methods* **42**(7), 1176–1192.

de Haan, L. and Resnick, S. I. (1977), 'Limit theory for multivariate sample extremes', *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **40**, 317–337.

Drees, H. and Kaufmann, E. (1998), 'Selecting the optimal sample fraction in univariate extreme value estimation', *Stochastic Processes and their Applications* **75**(2), 149–172.

Eastoe, E. F. and Tawn, J. A. (2009), 'Modelling non-stationary extremes with application to surface level ozone', *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **58**(1), 25–45.

Einmahl, J. H. J., de Haan, L. and Li, D. (2006), 'Weighted approximations of tail copula processes with application to testing the bivariate extreme value condition', *The Annals of Statistics* **34**, 1987–2014.

Einmahl, J. H. J., Li, J. and Liu, R. Y. (2009), 'Thresholding events of extreme in simultaneous monitoring of multiple risks', *Journal of the American Statistical Association* **104**, 982–992.

Elie, L. and Jeantheau, T. (1995), 'Consistency in heteroskedastic models', *Comptes Rendus de l'Académie des Sciences* **Serie I**, 1255–1258.

Engle, R. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Escobar-Bach, M., Goegebeur, Y. and Guillou, A. (2016), Local robust estimation of the Pickands dependence function. working paper or preprint.

Fisher, R. A. and Tippett, L. H. C. (1928), Limiting forms of the frequency distribution of the largest or smallest member of a sample, *in* 'Mathematical Proceedings of the Cambridge Philosophical Society', Vol. 24, Cambridge University Press, pp. 180–190.

Gudendorf, G. and Segers, J. (2012), 'Nonparametric estimation of multivariate extreme-value copulas', *Journal of*

Statistical Planning and Inference **142**, 3073–3085.

Hansen, P. R. and Lunde, A. (2005), 'A forecast comparison of volatility models: Does anything beat a garch(1,1)?', *Journal of Applied Econometrics* **20**, 873–889.

Hanson, T. E., de Carvalho, M. and Chen, Y. (2017), 'Bernstein polynomial angular densities of multivariate extreme value distributions', *Statistics and Probability Letters* .

Hastie, T. J. and Tibshirani, R. J. (1990), *Generalized Additive Models*, London: Chapman & Hall.

Heffernan, J. E. (2000), 'A directory of coefficients of tail dependence', *Extremes* **3**, 279–290.

Huang, X. (1992), Statistics of Bivariate Extreme Values, PhD thesis, Tinbergen Institute Research Series.

Huser, R., Davison, A. C. and Genton, M. G. (2016), 'Likelihood estimators for multivariate extremes', *Extremes* **19**, 79–103.

Knight, K. (2000), *Mathematical Statistics*, Chapman & Hall/CRC Press, Boca Raton.

Koijen, R. S. J., Philipson, T. J. and Uhlig, H. (2016), 'Financial health economics', *Econometrica* **84**, 195–242.

Kotz, S. and Nadarajah, S. (2000), *Extreme Value Distributions: Theory and Applications*, Imperial College Press.

Lamoureux, C. G. and Lastrapes, W. D. (1990), 'Heteroskedasticity in stock return data: Volume versus garch effects', *The Journal of Finance* **45**, 221–229.

Marcon, G., Padoan, S. A., Antoniano-Villalobos, I. et al. (2016), 'Bayesian inference for the extremal dependence', *Electronic Journal of Statistics* **10**(2), 3310–3337.

McNeil, A., Frey, R. and Embrechts, P. (2015), *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton University Press, Princeton, MA.

Mhalla, L., Chavez-Demoulin, V. and Naveau, P. (2017), 'Non-linear models for extremal dependence', *Available at SSRN: https://ssrn.com/abstract=2836587* .

Padoan, S. A., Ribatet, M. and Sisson, S. A. (2010), 'Likelihood-based inference for max-stable processes', *Journal of the American Statistical Association* **105**, 263–277.

Pauli, F. and Coles, S. (2001), 'Penalized likelihood inference in extreme value analyses', *Journal of Applied Statistics* **28**(5), 547–560.

Pickands, J. (1981), Multivariate extreme value distributions, *in* 'Proc. 43rd Session of the International Statistical Institute', pp. 859–878.

R Development Core Team (2016), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Resnick, S. I. (1987), *Extreme values, regular variation and point processes*, Springer, New York.

Sabourin, A. and Naveau, P. (2014), 'Bayesian dirichlet mixture model for multivariate extremes: A re-parametrization', *Computational Statistics & Data Analysis* **71**, 542–567.

Tsay, R. S. (2005), *Analysis of Financial Time Series*, Wiley, Hoboken.

Vatter, T. and Chavez-Demoulin, V. (2015), 'Generalized additive models for conditional dependence structures', *Journal of Multivariate Analysis* **141**, 147–167.

Wadsworth, J. L. and Tawn, J. A. (2013), 'A new representation for multivariate tail probabilities', *Bernoulli* **19**(5B), 2689–2714.

Wang, H. and Tsai, C.-L. (2009), 'Tail Index Regression', *Journal of the American Statistical Association* **104**(487), 1233–1240.

Xie, Y. (2015), *Dynamic Documents with R and Knitr*, Vol. 29, Chapman & Hall/CRC, Boca Raton, FL.

Yee, T. W. (2015), *Vector Generalized Linear and Additive Models: With an Implementation in R*, 1st edn, Springer, New York.

Yee, T. W. and Stephenson, A. G. (2007), 'Vector generalized linear and additive extreme value models', *Extremes* **10**(1-2), 1–19.

Yee, T. W. and Wild, C. J. (1996), 'Vector generalized additive models', *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **58**, 481–493.